Model reduction for slow–fast stochastic systems with metastable behaviour

Maria Bruna,[1, 2, a)] S. Jonathan Chapman,[1] and Matthew J. Smith[2]
[1)] *Mathematical Institute, University of Oxford, Oxford, OX2 6GG, U.K.*
[2)] *Computational Science Laboratory, Microsoft Research, Cambridge, CB1 2FB, U.K.*

(Dated: 23 April 2014)

The quasi-steady-state approximation (or stochastic averaging principle) is a useful tool in the study of multiscale stochastic systems, giving a practical method by which to reduce the number of degrees of freedom in a model. The method is extended here to slow–fast systems in which the fast variables exhibit metastable behaviour. The key parameter that determines the form of the reduced model is the ratio of the timescale for the switching of the fast variables between metastable states to the timescale for the evolution of the slow variables. The method is illustrated with two examples: one from biochemistry (a fast-species-mediated chemical switch coupled to a slower-varying species), and one from ecology (a predator–prey system). Numerical simulations of each model reduction are compared with those of the full system.

arXiv:1401.6601v3 [physics.chem-ph] 22 Apr 2014

I. INTRODUCTION

Understanding the impact of noise on nonlinear dynamical systems has been an active field of research for many years, with a wide range of applications in physics, chemistry, biology, ecology and earth science. Although the addition of noise sometimes does not change the qualitative dynamics and can be modelled by adding a stochastic perturbation to the deterministic solution trajectory (as in the linear noise approximation[1]), there is a growing number of applications in which noise has been shown to be crucial to explain features which cannot be captured by deterministic models. For example, in deterministic systems with multiple stable steady states the addition of even a small amount of noise causes these states to become metastable, with the stochastic system undergoing random transitions between the deterministic steady states. Examples include genetic regulatory networks,[2] lactose utilisation networks and the bet-hedging in bacteria.[3,4] This behaviour cannot be captured by the deterministic model,[5] but can be explained successfully with a stochastic model.[6]

A common feature of many complex dynamical systems is the presence of processes evolving on widely separated timescales. This can present a challenge for numerical simulation. Often we are only interested in the behaviour of solutions on a long timescale,[7] but in principle to determine this we need to resolve processes occurring on the fastest timescale. Metastability complicates this further, since the switching rate between states introduces a new implicit timescale to the process. For example, several ecological systems are believed to be able to switch between alternative states, such as between tree and grass-dominated vegetation.[8,9] A few key mechanisms may be sufficient to explain the observed bi-stability of these systems, but a wide range of other ecosystem components (such as the resident bird community) respond to these dynamics, and may do so at different rates; for example the population size of a resident bird community may respond relatively quickly to a sudden switch in the dominant vegetation type, but the carbon content of the soil may respond relatively slowly. This raises questions about how to efficiently and accurately simulate the dynamics of such systems: can the evolution of the slowly evolving components be predicted without simulating the full evolution of all the fast components?

In this paper we bring together these two features, multiple timescales and metastability, and consider noisy systems with slow and fast metastable components. We develop a method to remove the fast degrees of freedom while retaining the metastable behaviour in the resulting mathematical description of the slow processes.

Most of the work done to remove fast degrees of freedom from stochastic multiscale systems is based on extending the principle of slow or centre manifold theory of deterministic dynamical systems. This assumes that the phase space can be decomposed into fast variables x which relax very quickly and slow variables y, which change more slowly:

$$\epsilon \frac{\mathrm{d}x}{\mathrm{d}t} = F(x,y), \qquad x(0) = x_0, \tag{1a}$$

$$\frac{\mathrm{d}y}{\mathrm{d}t} = G(x,y), \qquad y(0) = y_0, \tag{1b}$$

where $\epsilon \ll 1$ quantifies the separation of timescales. The standard singular perturbation theory, based on Fenichel's theory,[10] consists of taking the limit $\epsilon \to 0$ and assuming that the fast variables x have equilibrated onto the *centre manifold*, $F(x,y) = 0$. The flow on this slow manifold (with reduced dimensionality) is then given by

$$\frac{\mathrm{d}y}{\mathrm{d}t} = G(h(y),y), \qquad y(0) = y_0, \tag{2}$$

where $F(h(y),y) = 0$.[11] Note that the projection onto the slow manifold is valid even if $F(x,y) = 0$ has more than one solution for y fixed; given the initial condition x_0, the fast system evolves deterministically to one of the steady states as determined by $F(x,y)$. However, this is

[a)] Electronic mail: bruna@maths.ox.ac.uk

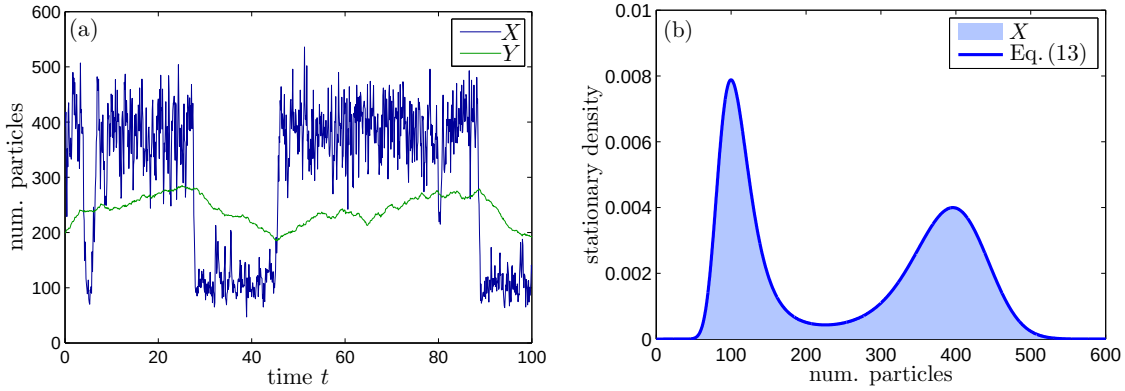

FIG. 1. (a) Single realisation of (6) obtained using the exact NRM SSA with $X(0) = 400$ and $Y(0) = 200$. (b) Comparison of the stationary distribution $\rho(x)$ given by (13) (solid dark blue line) and the histogram of X from SSA of (6) (10^7 samples used, light blue area). We use $\tau_y = 25$ and the parameter values (7).

no longer true when we introduce noise in the system: noise can make $x(t)$ fluctuate between several states satisfying $F(x, y) = 0$ and (2) does no longer capture the dynamics of $y(t)$. Developing a good approximation for the dynamics of system (1) in the presence of multiple stable states and noise is the key problem that we address here.

There exists a variety of approximation methods to generalise this model reduction to stochastic systems. The theoretical foundation of these can be traced back to the stochastic centre manifold theory developed in Ref. 12. A good introduction to the existing techniques, classified according to the stochastic system representation, can be found in Ref. 13. Consider the following stochastic counterpart of (1):

$$\mathrm{d}X = \frac{1}{\epsilon}F(X,Y)\mathrm{d}t + \frac{1}{\sqrt{\epsilon}}\sigma(X,Y)\mathrm{d}W_X(t),\ X(0) = X_0, \tag{3a}$$

$$\mathrm{d}Y = G(X,Y)\mathrm{d}t + \Sigma(X,Y)\mathrm{d}W_Y(t), \quad Y(0) = Y_0, \tag{3b}$$

where $W_X(t)$ and $W_Y(t)$ are independent standard Brownian motions. Even though the presence of noise makes rigorous analysis more complicated, the idea is conceptually very simple: one assumes that, "freezing" the slow variables at $Y(t) = \hat{y}$, the fast process $X(t)$ reaches a unique stationary density $\rho^{\hat{y}}(x)$ (or rather a *quasi-steady density*), which is the analogue of the unique steady-state in the deterministic system. The slow stochastic system analogous to (2) is then obtained by averaging in the fast variables over this density:

$$\mathrm{d}Y = \overline{G}(Y)\mathrm{d}t + \overline{\Sigma}(Y)\mathrm{d}W_Y(t), \tag{4}$$

where

$$\overline{G}(Y) = \int G(x,Y)\rho^{\hat{y}}(x)\,\mathrm{d}x, \quad \overline{\Sigma}(Y) = \int \Sigma(x,Y)\rho^{\hat{y}}(x)\,\mathrm{d}x, \tag{5}$$

a process known as *stochastic averaging*.[14] This process is sometimes referred to as the quasi-steady-state (QSS) reduction.[15] The changes in $Y(t)$ may push $X(t)$ away from its previous quasi-steady density, but by assumption fast transients in $X(t)$ die out quickly and a new quasi-steady density $\rho^{\hat{y}}(x)$ applies. Numerical schemes based on the QSS are analysed in Ref. 7 and implementations of those can be found in Ref. 16 and 17.

The underlying assumption of the above method is that $X(t)$ is well approximated by a random variable chosen from its steady-state distribution $\rho^{\hat{y}}(x)$, conditioned on a fixed value of $Y(t) = \hat{y}$.[18] However, this might not be true if the fast degrees of freedom exhibit metastability. Our goal in this paper is to extend the model reduction of slow–fast stochastic systems to scenarios in which the fast variables do not have a single invariant measure over the timescales of interest, but switch randomly between a number of invariant measures.

Consider the stochastic metastable process for a single variable $X(t)$. The standard example of metastable behaviour is Brownian motion with a double well potential [resulting in a stationary density similar to the one shown in Fig. 1(b)]. On short timescales the particle is most likely found near one of the two potential minima, but on long timescales the particle can transition over the energy barrier that separates the two wells. This problem itself can be thought of a slow–fast system, where the fluctuations of $x(t)$ within a well is the fast process, and the metastable transitions between wells play the role of a (discrete) slow process, since they typically occur on a timescale $\tau_s \gg \epsilon$.

The aim of model reduction techniques for metastable processes is to eliminate the fast degrees of freedom while retaining the slow metastable transitions. The extension of the QSS to metastability is known as quasi-stationary analysis (QSA) and is based on WKB projection methods.[19] The main feature of model reduction in this context is that, instead of averaging over a global quasi-stationary density $\rho(x)$ (as in the QSS, which would average over all the metastable basins), it is assumed that $X(t)$ relaxes in one basin of attraction B_j

and is well-approximated by a stationary random process chosen from the quasi-stationary density $\rho_j(x)$ restricted on basin B_j. Then the reduced low-dimensional model consists of a discrete jump process between states j, where the switching rates are calculated from the transition rates of the original process $X(t)$.[18,20,21] In higher-dimensional systems where analytical approximations are not possible, computational techniques can be applied to determine the state space of the metastable variables and to sample the transition probabilities. A recent survey of such techniques applied to the simulation of rare events in molecular dynamics (such as conformation changes) can be found in Ref. 22.

Now suppose X is coupled to an additional slow variable Y. For ease of exposition we focus on the simpler case in which the dynamics of X are independent of Y. As we will discuss in section V, the extension to the more general case is conceptually straightforward. We see that three timescales are required to characterise such a system: the timescale of the fast variables ϵ, the timescale of the *switches* of the fast variables $\tau_s \gg \epsilon$, and the timescale of the slow variables $\tau_y \gg \epsilon$. We are interested in approximating the system on the slowest of these timescales.

We will examine how the reduced model changes with the relative size of τ_s and τ_y. If $\tau_y \ll \tau_s$, the reduced model is simply a discrete jump process.[20] If $\tau_y \gg \tau_s$, the slow variables evolve on a timescale which is even slower than the switches in x and the classical QSS approximation applies. In the intermediate regime in which $\tau_y = \mathcal{O}(\tau_s)$ things become more interesting, as the evolution of y critically depends on the metastable behaviour of x. This is the main contribution of the present work.

The rest of this paper is organised as follows. In Sec. II we introduce a model chemical system with which we will illustrate the model reduction method. In section III we describe the idea behind our analysis and show how to obtain a reduced model in the slow timescale when the fast process is monostable, and how this changes when there is metastability. We introduce the three characteristic timescales in our slow–fast model system with bistability, from which we define three regimes depending on their relative sizes. We present the three corresponding reduced models and show numerical comparisons between these reduced slow-scale stochastic models and the full slow–fast process using a stochastic simulation algorithm (SSA). In section IV we apply the method to a simple predator–prey model with bistable prey and use it to estimate the mean extinction time of the predator. More background on why we choose this particular case study will be given in the beginning of that section. In the concluding section V we briefly discuss the results in the context of current directions of research in theoretical ecology and stochastic nonlinear dynamics. The justification of the reduced models through a perturbation analysis of the corresponding Fokker–Planck equations is presented in appendix A.

II. MODEL PROBLEM

We introduce the following model as an example of a slow–fast system with bistability. A system with two chemical species X and Y changes according to the reactions

$$\emptyset \underset{k_2/\epsilon}{\overset{k_1/\epsilon}{\rightleftharpoons}} X, \qquad 2X \underset{k_4/\epsilon}{\overset{k_3/\epsilon}{\rightleftharpoons}} 3X, \tag{6a}$$

$$\emptyset \underset{k_6/\tau_y}{\overset{k_5(X)/\tau_y}{\rightleftharpoons}} Y \tag{6b}$$

Note that the evolution of the chemical X is self-determining, while the production of chemical Y depends on the amount of X present in the system through the function $k_5(x)$. We suppose the system is in a well-mixed reactor of unit volume. We have chosen to write the reaction rates for $X(t)$ in terms of fixed k_1 to k_4 and a small parameter ϵ; by varying ϵ we can change the timescale for the evolution of X without changing the equilibrium distribution. Similarly, we have chosen to write the reaction rates for $Y(t)$ in terms of fixed $k_5(x)$ and k_6 and a parameter τ_y.

The set of four reactions (6a) governing $X(t)$ is known as the Schlögl model.[23] This model has two favourable states (or basins of attraction) for some parameter values, and the stochastic system displays bistable switching between these states even for parameter values for which the deterministic model is monostable.[24]

The stochastic model of the six chemical reactions (6) introduces the propensity functions

$$\begin{aligned}
&\alpha_1(x) = k_1/\epsilon, && \alpha_2(x) = k_2 x/\epsilon,\\
&\alpha_3(x) = k_3 x(x-1)/\epsilon, && \alpha_4(x) = k_4 x(x-1)(x-2)/\epsilon,\\
&\alpha_5(x) = k_5(x)/\tau_y, && \alpha_6(y) = k_6 y/\tau_y,
\end{aligned}$$

which are such that the probability of reaction i happening in time $\mathrm{d}t$ is $\alpha_i \mathrm{d}t$.[25] The system may then be simulated by using a stochastic simulation algorithm (SSA) such as the exact Gillespie's Direct Method.[26] Here we use the equivalent exact and efficient Next Reaction Method (NRM).[27]

The timescale for the evolution of X is set by the decay rate k_2/ϵ. Similarly the timescale for the evolution of Y is set by the decay rate k_6/τ_y. Thus if ϵ and τ_y are to represent these timescales we should choose k_2 and k_6 to be of order unity. The other rates determine the equilibrium values of X and Y. If we suppose a typical equilibrium value of X is $O(\delta^{-1})$, this corresponds to $k_1 = O(\delta^{-1})$, $k_3 = O(\delta)$, $k_4 = O(\delta^2)$. In Fig. 1(a) we plot one realisation of the SSA of (6) obtained for $X(0) = 400$, $Y(0) = 200$, $\tau_y = 25$, and

$$\begin{aligned}
&k_1 = 0.6\,\delta^{-1}, && k_2 = 1, && k_3 = 0.48\,\delta,\\
&k_4 = 0.0666\,\delta^2, && k_5(x) = 0.833\,x, && k_6 = 1,\\
&\epsilon = 0.025, && \delta = 0.01.
\end{aligned} \tag{7}$$

The parameters values are chosen such that system (6) is bistable in $X(t)$ (as in Ref. 24) and $Y(t)$ evolves in

the same timescale as switches in X. The rate $k_5(x)$ (which we will also use in the simulations that follow) is such that the production reaction for Y can be written as $X \longrightarrow X + Y$ with rate $0.833/\tau_y$. We also plot the stationary distribution of X, obtained by a long time simulation of the SSA with final time $t = 10^6$ and recording the value of X every $\Delta t = 0.1$ in Fig. 1(b). We see that for these parameter values the system (6) has two favourable states for X. We see that $X(t)$ spends quite a long time fluctuating around one of these states before switching to the other. In section A 2 we will derive an analytical approximation of the mean switching times between states and find that is exponentially large in the parameter δ. We want to understand how the oscillation between these two states influences the dynamics of Y, and obtain an efficient and accurate way to take this effect into account without having to simulate the full X dynamics.

A. Continuous approximation: the chemical Langevin and the Fokker–Planck equations

For ease of exposition we will apply our model reduction techniques not to (6) directly, but to the continuous approximation of this system given by the chemical Langevin equation. This approximation is similar to the numerical method of τ-leaping.[28] The idea of the approximation is that for large X and Y the propensity functions do not change significantly after each individual reaction event, which enables us to jump forward many reaction events without updating the propensities. The change in the number of molecules between timesteps is then Poisson distributed. If we now take a limit in which many reactions are considered between timesteps, but still with a small relative change in the number of molecules, X and Y may be approximated by continuous random variables and the change in molecular numbers between timesteps is approximately normally distributed (see Refs. 24 and 29 for more details of this approximation). In that case we arrive at the chemical Langevin equations

$$\mathrm{d}X = \frac{v(X)}{\epsilon}\mathrm{d}t + \sqrt{\frac{2d(X)}{\epsilon}}\,\mathrm{d}W_X(t), \tag{8a}$$

$$\mathrm{d}Y = \frac{V(X,Y)}{\tau_y}\mathrm{d}t + \sqrt{\frac{2D(X,Y)}{\tau_y}}\,\mathrm{d}W_Y(t), \tag{8b}$$

which are of the same form as the slow–fast system of stochastic differential equations (3). The drift coefficients $v(x)$, $V(x,y)$ and the diffusion coefficients $d(x)$ and $D(x,y)$ are given by

$$\begin{aligned} v(x) &= k_1 - k_2 x + k_3 x(x-1) - k_4 x(x-1)(x-2), \\ d(x) &= \frac{1}{2}[k_1 + k_2 x + k_3 x(x-1) + k_4 x(x-1)(x-2)], \\ V(x,y) &= k_5(x) - k_6 y, \\ D(x,y) &= \frac{1}{2}[k_5(x) + k_6 y]. \end{aligned} \tag{9}$$

Let $P(x,y,t)\mathrm{d}x\mathrm{d}y$ be the probability that $X(t) \in [x, x+\mathrm{d}x)$ and $Y(t) \in [y, y+\mathrm{d}y)$ at time t. Then the chemical Fokker–Planck (FP) equation for the joint probability distribution function corresponding to (8) is given by

$$\begin{aligned} \frac{\partial P}{\partial t}(x,y,t) =& \frac{1}{\epsilon}\frac{\partial}{\partial x}\left\{\frac{\partial}{\partial x}[d(x)P] - v(x)P\right\} \\ &+ \frac{1}{\tau_y}\frac{\partial}{\partial y}\left\{\frac{\partial}{\partial y}[D(x,y)P] - V(x,y)P\right\}. \end{aligned} \tag{10}$$

This equation is complemented with no-flux boundary conditions on $[0,\infty)^2$ since the probability P must remain normalised for all times.

The stationary distributions of X and Y, $\rho(x)$ and $q(y)$ respectively, correspond to taking the limit $t \to \infty$ in the marginal probability densities

$$\rho(x) = \lim_{t\to\infty}\int P(x,y,t)\,\mathrm{d}y, \quad q(y) = \lim_{t\to\infty}\int P(x,y,t)\,\mathrm{d}x. \tag{11}$$

Since the dynamics of X do not depend on Y, it is straightforward to solve for $\rho(x)$ by integrating (10) with respect to y to give

$$\frac{\partial}{\partial x}\left(\frac{\partial}{\partial x}[d(x)\rho] - v(x)\rho\right) = 0, \tag{12}$$

where we have used no-flux boundary conditions at $y = 0$ and $y = \infty$. Thus the stationary distribution of X is given by

$$\rho(x) = \frac{A}{d(x)}\exp\left[\int_0^x \frac{v(s)}{d(s)}\mathrm{d}s\right], \tag{13}$$

where A is the normalisation constant. The solution (13) is plotted in Fig. 1(b) as a blue solid line. As expected, the results compare well with the results obtained by the long time stochastic simulations.

III. REDUCED SLOW-SCALE MODELS VIA STOCHASTIC AVERAGING

In this section we show how the stochastic model (6) may be reduced when ϵ is small, and how the reduced model depends on the size of τ_y relative to the size of τ_s.

A. The case when $X(t)$ is monostable

To give some context to the reduced models which follow, we first consider the simpler case in which the fast species X is monostable. In Fig. 2(a) we present a stochastic simulation of (6) with $X(0) = Y(0) = 0$ for $\delta = 0.01$, $\epsilon = 0.025$, $\tau_y = 25$, $k_1 = 0.8\delta^{-1}$, and k_2 to k_6 as in (7). In contrast to Fig. 1, we see that for this value of k_1 the stochastic model fluctuates about a single value. This is demonstrated in Fig. 2(b), where we plot

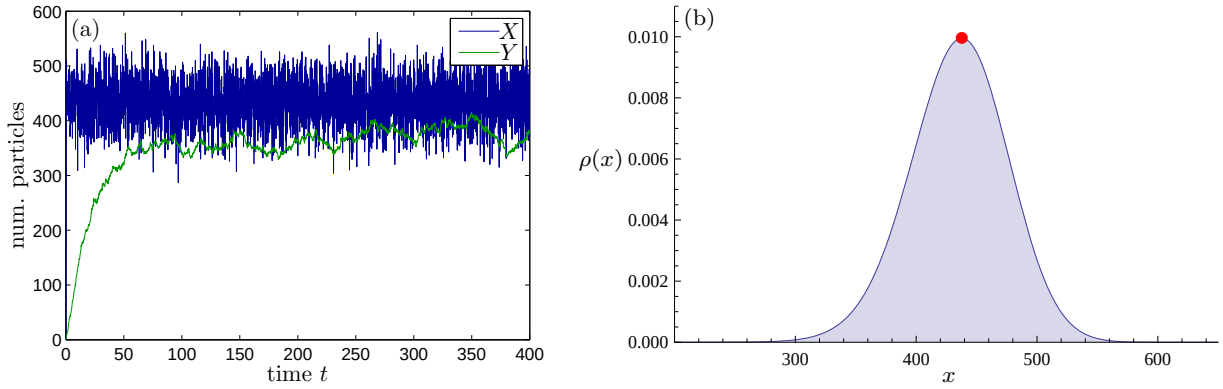

FIG. 2. (a) Single realisation of (6) obtained using the exact NRM SSA. (b) Marginal stationary distribution $\rho(x)$ given by (13) with the stable point x_0 marked with the red circle. We use $\tau_y = 25$, $X(0) = Y(0) = 0$, and the parameter values (7) except that $k_1 = 0.8\delta^{-1}$.

the stationary distribution $\rho(x)$ using (13). The maximum of the stationary distribution is $x_0 \approx 437.8$, which satisfies $v(x_0) - d'(x_0) = 0$.

In this case, since the relaxation time of X to its equilibrium density is much faster than the decay time of the slow reaction, the slow chemical reactions (6b) can be well-approximated by[16,17]

$$\emptyset \underset{k_6/\tau_y}{\overset{\overline{k}_5/\tau_y}{\rightleftharpoons}} Y \tag{14}$$

where $\overline{k}_5$ is the average production rate, given by

$$\overline{k}_5 = \int_0^\infty k_5(x)\rho(x)\,\mathrm{d}x. \tag{15}$$

For our particular example in which $k_5(x) = k_5 x$, the effective production rate is simply $\overline{k}_5 = k_5\overline{X}$.

The analogous quasi-steady-state reduction in the chemical Langevin equation (8b) is

$$\mathrm{d}Y = \frac{\overline{V}(Y)}{\tau_y}\,\mathrm{d}t + \sqrt{\frac{2\overline{D}(Y)}{\tau_y}}\,\mathrm{d}W_Y(t),$$

where

$$\overline{V}(y) = \int V(x,y)\rho(x)\,\mathrm{d}x, \qquad \overline{D}(y) = \int D(x,y)\rho(x)\,\mathrm{d}x.$$

like in the stochastic averaged model (4) and (5) described in Sec. I.

B. The case when $X(t)$ is bistable

Let us now consider the case when $X(t)$ is bistable, that is, it switches between two favourable states as shown in Fig. 1. Figure 3(a) shows an illustrative trajectory of the system (6) for $X(0) = Y(0) = 0$, with parameter values given by (7). Figure 3(b) shows the stationary density $\rho(x)$ computed using (13). We denote the two peaks of this density by $x_\pm$ and the relative minimum (or unstable node) as x_*. Now it is not clear whether it is appropriate to use the standard stochastic averaging technique, since the assumption that $X(t)$ converges quickly to a stationary process with measure $\rho(x)$

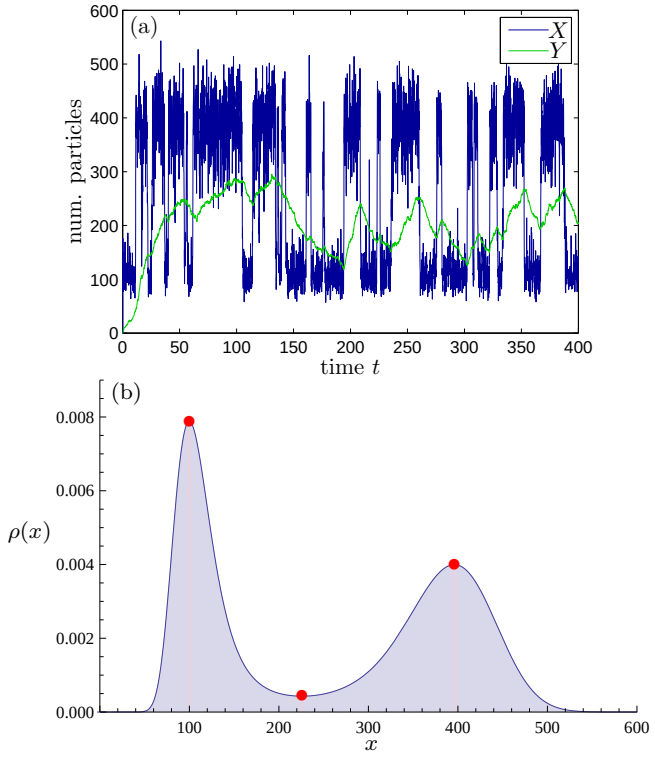

FIG. 3. (a) Single realisation of (6) obtained using the exact NRM SSA. (b) Marginal stationary distribution $\rho(x)$ given by (13) with the two stable fixed points $x_\pm$ and unstable fixed point x_* marked with red circles. We use $\tau_y = 25$, $X(0) = Y(0) = 0$, and the parameter values (7).

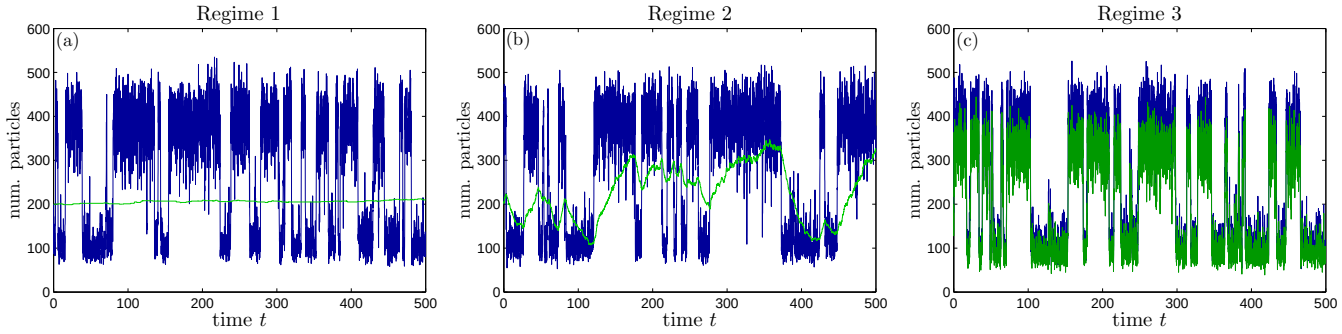

FIG. 4. Evolution of $X(t)$ (blue lines) and $Y(t)$ (green lines) from a single realisation of (6) obtained using the exact SSA. We use the parameter values (7), $X(0) = Y(0) = 200$, and vary τ_y to change between regimes. (a) Regime 1: $\tau_y = 2500$. (b) Regime 2: $\tau_y = 25$. (c) Regime 3: $\tau_y = 0.25$.

is challenged by the metastable behaviour of $X(t)$.

To obtain a reduced slow system now we first need to characterise this metastable behaviour. In particular, we see that a third timescale τ_s emerges in the problem as the characteristic time for switches of the fast variable $X(t)$ between the two favourable states. Typically this timescale is much longer than the relaxation timescale for $X(t)$ within each well (indeed, this is the definition of metastability used here).

Thus we have the following scenario. On a short timescale $X(t)$ relaxes to a quasi-stationary distribution centred around one of x_- or x_+. On the longer timescale of τ_s the system switches from one of these distributions to the other, as $X(t)$ makes the transition between wells. The nature of the reduced mode depends crucially on the relative sizes of the switching timescale τ_s and the timescale for the slow process τ_y. We will see that there are three parameter regimes, corresponding to $\tau_s \ll \tau_y$, $\tau_s \sim \tau_y$ and $\tau_s \gg \tau_y$, respectively. Illustrative simulations of each of these regimes are shown in Fig. 4. In Subsections III C–III E we present the reduced stochastic model appropriate for each regime, together stochastic simulations comparing the reduced model to the full system. The mathematical justification for the reduced models is given in appendix A.

C. Regime 1: $\epsilon \ll \tau_s \ll \tau_y$

A sample trajectory in this regime is shown in Figure 4(a). In this case, X has time to fully equilibrate on the timescale of the evolution of Y. Thus the standard stochastic averaging can be used, and the effective production rate of Y can be computed as if X was monostable as

$$\emptyset \underset{k_6/\tau_y}{\overset{\overline{k}_5/\tau_y}{\rightleftharpoons}} Y, \qquad \overline{k}_5 = \int_0^\infty k_5(x)\rho(x)\,\mathrm{d}x. \tag{16}$$

We call (16) *reduced model 1* or RM1. A derivation of this model using the associated Fokker–Planck equation is given in A 1.

In Figures 5 and 6 we compare simulation results (using the NRM SSA) of the full system (6) and the reduced system RM1. We use a long timescale for $Y(t)$, $\tau_y = 2500$, so that the system is in Regime 1. The stationary marginal density $q(y)$ (histograms obtained by a long time simulation of both models) is plotted in Fig. 5. Clearly, the reduced model provides a very good approximation to the histogram coming from the full-system simulation. The benefit of eliminating the fast variable is illustrated by the fact that the RM1 histogram took only 0.8 seconds to compute while the full system histogram required over 14 hours of computing time (using a standard desktop computer). In Fig. 6 we compare the time-dependent behaviour of the exact full model and RM1, plotting the mean and standard deviation of $Y(t)$, μ_Y and σ_Y respectively. We initialise the system with $X(0) \sim \rho(x)$ and $Y(0) = 0$. While the mean obtained

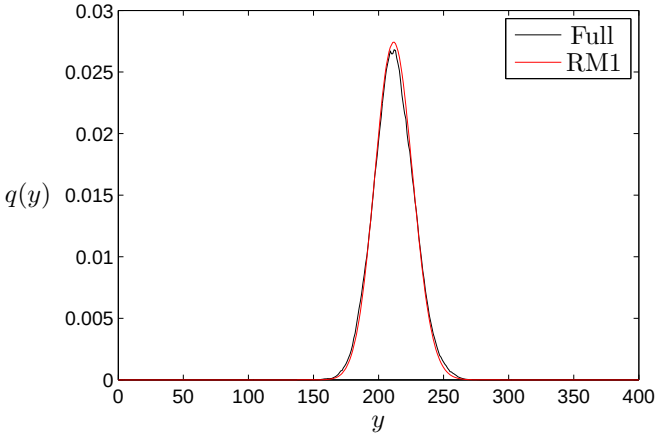

FIG. 5. Comparison of the marginal stationary density $q(y)$ in regime 1: full system (6) (solid black line) and approximate reduced model 1 (16) (dash red line). Both histograms are computed by running the NRM algorithm up to $t = 10^7$ taking recordings every $\Delta t = 5$ (2×10^6 samples). We use $\tau_y = 2500$ and the parameter values (7).

from the RM1 approximation is indistinguishable from that of the full system, the standard deviation is slightly underestimated by RM1. We note that the error in σ_Y in Fig. 6(b) is in fact quite small relative to the mean μ_Y (0.65%), and this is why this difference is only just recognisable in Fig. 5.

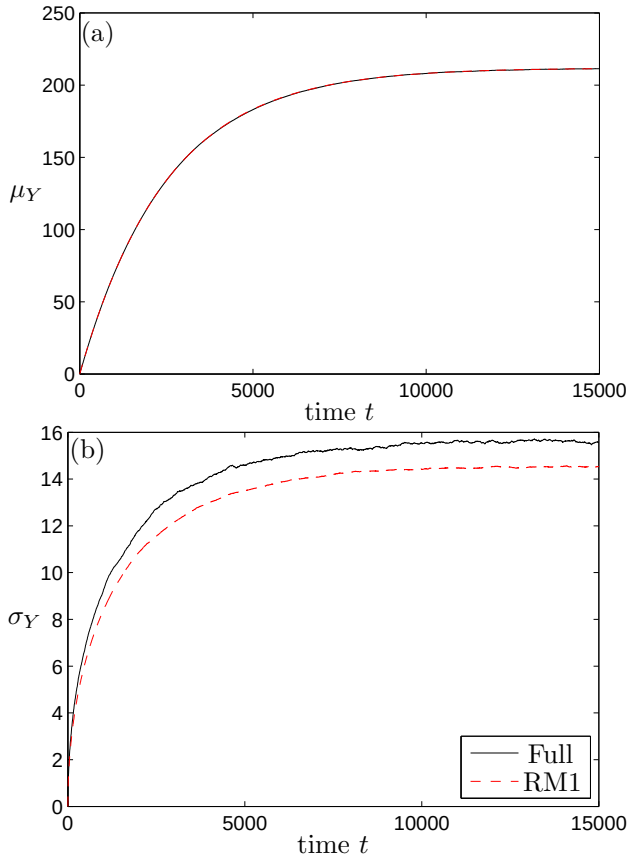

FIG. 6. Comparison of the time-dependent (a) mean $\mu_Y(t)$ and (b) standard deviation $\sigma_Y(t)$: full system (6) (solid black line) and approximate reduced model 1 (16) (dash red line). Black and red curves are computed as the average over 1.1×10^4 and 10^5 realisations respectively with initial conditions $Y(0) = 0$ and $X(0) \sim \rho(x)$ (for (6) only). We use $\tau_y = 2500$ and the parameter values in (7).

D. Regime 2: $\epsilon \ll \tau_s \sim \tau_y$

When the timescale for the evolution of species Y is of the same order as that of the switches in $X(t)$, it is important to keep the bistable nature of the system in the reduced model, since Y responds differently depending on which well X is in. Fig. 4(b) shows an example of the system in this regime. As in the one-dimensional example described in Sec. I, using a QSA the bistability can be kept in by introducing a discrete two state stochastic process $S(t)$ governing the jumps of $X(t)$: $S(t) = S_-$ when $X(t)$ is in the left well $\Omega_- = [0, x_*)$ and $S(t) = S_+$ when $X(t)$ is in the right well $\Omega_+ = (x_*, \infty)$. Within each well X quickly reaches a quasi-equilibrium, and therefore we can use a modified stochastic averaging conditional on the value of S. Then the reduced stochastic system is

$$S_- \underset{k_+}{\overset{k_-}{\rightleftharpoons}} S_+, \qquad \emptyset \underset{k_6/\tau_y}{\overset{\overline{k}_5(S)/\tau_y}{\rightleftharpoons}} Y \tag{17a}$$

where the effective production rate takes one of two values depending on whether X is in the left or right well,

$$\overline{k}_5(S_\pm) = \int_{\Omega_\pm} k_5(x)\rho_\pm(x)\,\mathrm{d}x; \tag{17b}$$

here $\rho_\pm$ is the normalised stationary density of X conditional on being in the left (minus) or right (plus) well. This definition corresponds to imposing a reflecting boundary condition on x_* in (12).[30] A formal derivation of this model is given in A 2. We denote this reduced system by *reduced model 2* or RM2. The rate constants $k_\pm$ of the process for $S(t)$ are the inverse of the mean switching times $T_\pm$ for X to jump from one well to the other [the switching timescale τ_s is such that $T_\pm \sim O(\tau_s)$]. Determining these rates accurately is one of the main challenges of Regime 2; we will return to this issue in section III F. For the chemical $X(t)$, we will find that

$$k_- = 0.0931, \qquad k_+ = 0.0835. \tag{17c}$$

A comparison of a sample trajectory of the system obtained from a NRM simulation of the full system and of the reduced system RM2 is shown in Fig. 7. We choose the timescale of $Y(t)$ such that $\tau_y \sim \tau_s$ and the system is in Regime 2. In Fig. 7(a), the $X(t)$ and $Y(t)$ populations simulated from the full model (6) are shown in dark blue and green solid lines, respectively. The two-valued light blue curve represents the bistable switching of $X(t)$ and takes as values the conditional average values

$$\overline{X}_\pm = \int_{\Omega_\pm} x\rho_\pm(x)\,\mathrm{d}x, \tag{18}$$

according to when $X(t) \in \Omega_\pm$. Fig. 7(b) shows one trajectory of the reduced model (17). In RM2, $X(t)$ is not explicitly simulated, but we can still illustrate an approximate trajectory using the switch variable $S(t)$: in the first half of Fig. 7(b), we plot $\overline{X}_\pm$ as specified by $S(t) = S_\pm$ as well as the corresponding 95% confidence interval of $\rho_\pm(x)$ (dash blue lines). In the second half, to allow for an easier comparison with Fig. 7(a), we plot instead samples X from $\rho_\pm(x)$ depending on the value of $S(t)$ (taken every $\Delta t = 0.1$). Using this procedure, we can make a run of the RM2 look very similar to a full model run as in Fig. 7(a).

While Fig. 7 indicates that the qualitative behaviour of the reduced model (17) is similar to that of the full system, a more quantitative comparison is appropriate. In Fig. 8 we compare the stationary distributions of Y (histograms) obtained from a long-time NRM simulation

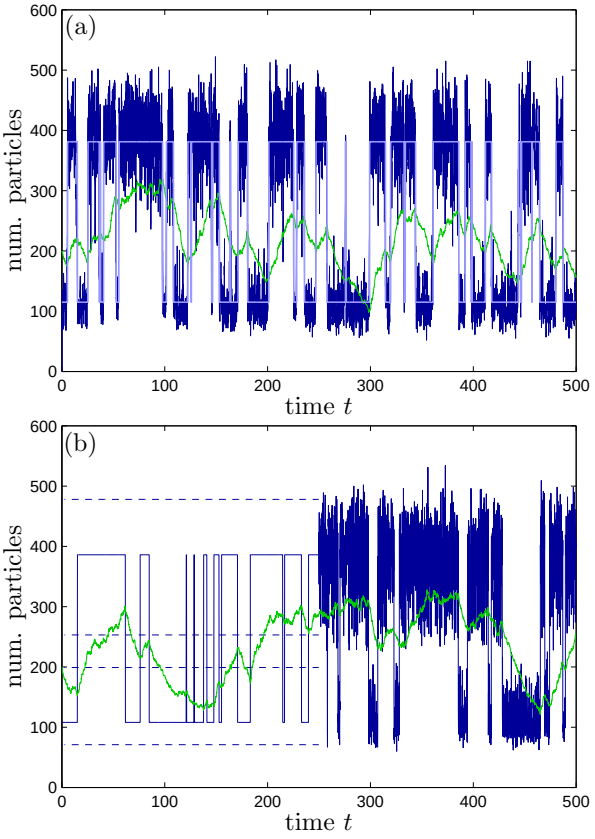

FIG. 7. Comparison of one single realisation in Regime 2 of (a) the exact full system (6) and (b) the approximated reduced model 2 (17) using the NRM algorithm. We use $\tau_y = 25$, $X(0) = 0$ [*i.e.* $S(0) = S_-$ in (b)], $Y(0) = 200$, and the parameter values (7). In (a), the simulated quantities $X(t)$ and $Y(t)$ are shown in dark blue and green, respectively. The lighter blue line shows reduced two-valued curve $\overline{X}_\pm$ corresponding to the bistable switch [$\overline{X}_\pm$ if $X(t) \in \Omega_\pm$, see (18)]. In (b), the simulated variable $Y(t)$ is shown in green, and the approximate representation of $X(t)$ from the simulated $S(t)$ in shown as follows: in the first half, $\overline{X}_\pm$ (according to $S(t) = S_\pm$) and the 95% confidence intervals of $\rho_\pm(x)$ are plot as solid and dash blue lines, respectively. In the second half, the solid blue line shows samples taken every $\Delta t = 0.1$ from $\rho_\pm(x)$ as per $S(t) = S_\pm$.

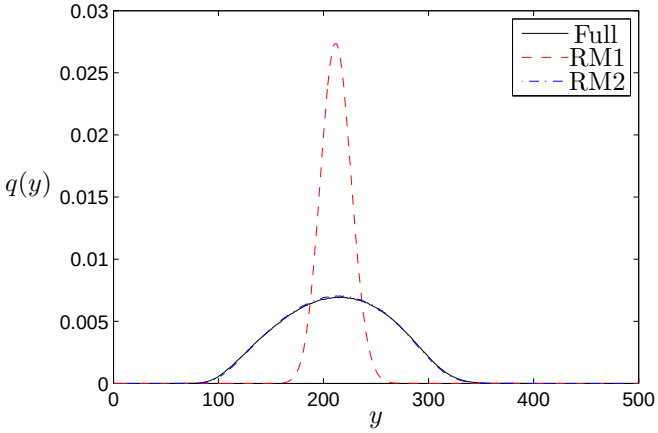

FIG. 8. Comparison of the marginal stationary density $q(y)$ in regime 2: full system (6) (solid black line) and approximate reduced models 1 (16) (dash red line) and 2 (17) (dot-dash blue line). All three histograms are computed by running the NRM algorithm up to $t = 10^7$ taking recordings every $\Delta t = 0.5$ (2×10^7 samples). We use $\tau_y = 25$ and the parameter values (7).

of the full model (6) and the reduced model 2 (17). To illustrate the need for a new model reduction in Regime 2, we also plot the histogram obtained by a standard stochastic averaging using RM1 (16). The RM2 solution is in excellent agreement with the full system, while the RM1 approximation considerably underestimates the variance of the distribution. In Fig. 9 we show a comparison of the mean and standard deviation of $Y(t)$ for the three models. We initialise the system with $Y(0) = 0$ and $X(0) \sim \rho(x)$. Note that the latter is not required for RM1 since $X(t)$ has been eliminated from that model; in contrast, the initialisation for RM2 is $S(0) = S_\pm$ depending on the well in which the sample from $\rho(x)$ is in [namely, $S(0) = S_-$ with probability $\int_{\Omega_-} \rho(x)\mathrm{d}x$, and S_+ otherwise]. We see that while the RM1 (16) captures the behaviour of the mean of Y fairly well, it significantly underestimates the variance of Y.

E. Regime 3: $\epsilon \sim \tau_y \ll \tau_s$

Now consider a scenario in which $\tau_y \ll \tau_s$, as illustrated in Fig. 4(c). Species Y is now fast by comparison to the switches of X between wells. Thus on the slowest timescale of τ_s both X and Y have time to reach a quasi-equilibrium, conditional on X being in a given well. In this case, when we eliminate the fast variables we are left only with the binary switch $S(t)$. Hence the reduced system in Regime 3 on the τ_s timescale is simply

$$S_- \underset{k_+}{\overset{k_-}{\rightleftharpoons}} S_+, \tag{19}$$

where $k_\pm$ are given in (17c); see A 3 for more details. We call this the *reduced model 3* (RM3). When $S(t) = S_-$ (resp. S_+), the system has a quasi-stationary density $Q_-(x, y)$ [resp. $Q_+(x, y)$], which is the stationary density conditional on X being in the left (resp. right) well. To clarify this, in Fig. 10 we plot the joint stationary density $Q(x, y)$ and the two marginal stationary densities $\rho(x)$ and $q(y)$ obtained by the long-time exact simulation of the full system (6) with τ_y such that the system is in Regime 3. This confirms that $Y(t)$ is also bimodal in this regime [even though the switches between its two modes are still controlled by $X(t)$]. To obtain the quasi-stationary densities $Q_\pm(x, y)$ [and the corresponding marginals $\rho_\pm(x)$ and $q_\pm(y)$], we follow a similar pro-

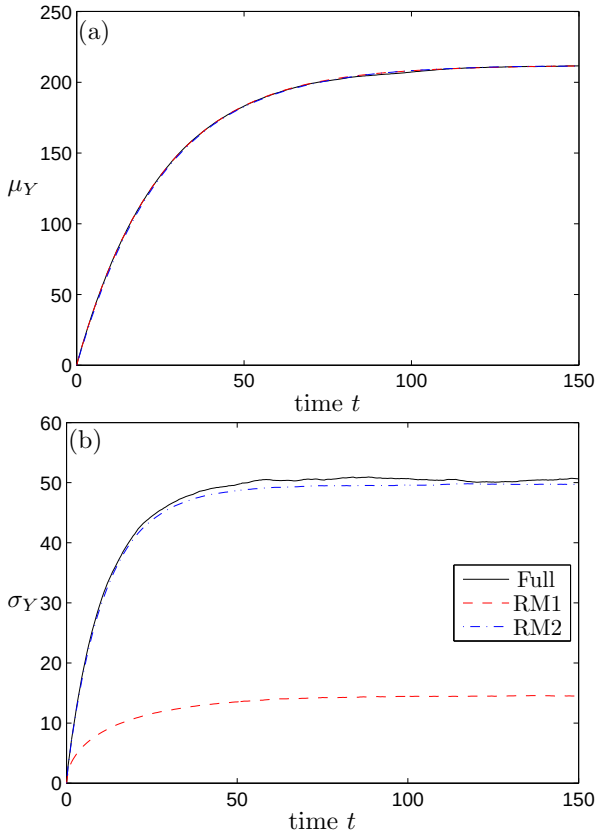

FIG. 9. Comparison of the time-dependent (a) mean $\mu_Y(t)$ and (b) standard deviation $\sigma_Y(t)$: full system (6) (solid black line) and approximate reduced models 1 (16) (dash red line) and 2 (17) (dot-dash blue line). Black, red, and blue curves are computed as the average over 1.1×10^4, 10^5, and 10^5 realisations respectively with initial conditions $Y(0) = 0$ and $X(0) \sim \rho(x)$. We use $\tau_y = 25$ and the parameter values in (7).

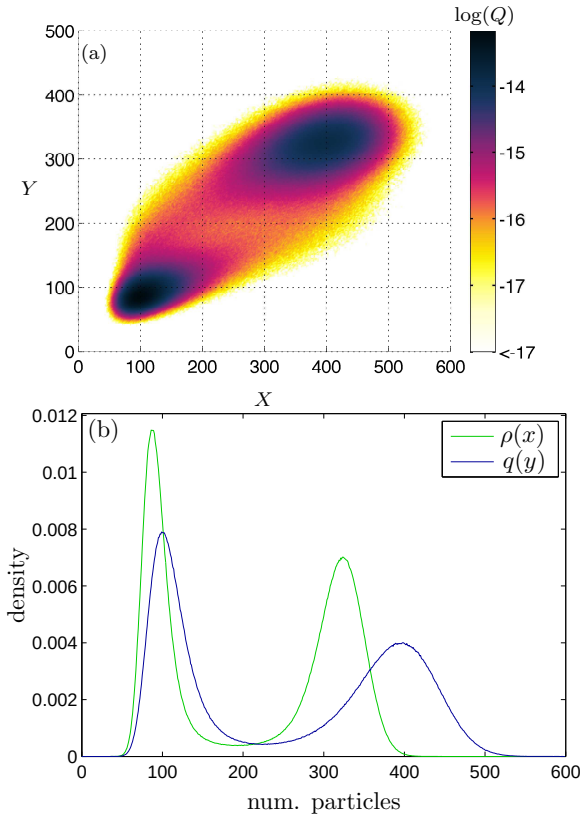

FIG. 10. (a) Logarithm of the joint stationary density $Q(x, y)$ and (b) marginal stationary densities $\rho(x)$ and $q(y)$ measured from a long simulation (up to $t = 10^6$ taking recordings every $\Delta t = 0.1$) of the full system (6) (blue and green histograms, respectively). We use $\tau_y = 0.25$ and the parameters values (7).

ceeding using a long-time simulation of (6) but with a reflecting boundary condition at x_*.

In Fig. 11 we compare a sample trajectory of the full set of reactions (6) with sample trajectory of the reduced system (19). We choose $\tau_y = 0.25$ so that the system is in Regime 3. Figures 11(a) and (c) show the evolution of $X(t)$ and $Y(t)$ respectively from (6), together with the two-valued filtered curves to represent the switches in $X(t)$ given by

$$(X(t), Y(t)) = \begin{cases} (\overline{X}_-, \overline{Y}_-), & \text{if } S(t) = S_-, \ (X(t) < x_*), \\ (\overline{X}_+, \overline{Y}_+), & \text{if } S(t) = S_+, \ (X(t) > x_*), \end{cases} \tag{20}$$

where

$$\overline{X}_\pm = \int_{\Omega_\pm} x Q_\pm(x, y)\, \mathrm{d}x\mathrm{d}y, \tag{21}$$

and similarly for $\overline{Y}_\pm$. Fig. 11(b) and (d) show one trajectory of the RM3 (19), where the output $S(t)$ is mapped to $X(t)$ and $Y(t)$ units using the quasi-stationary densities $\rho_\pm(x)$ and $q_\pm(y)$ [similarly to what we did in Fig. 7(b)]. In the first half of Figs. 11(b,d) we plot $(X(t), Y(t))$ (20) (solid lines) and the 95% confidence intervals (dashed lines) of $\rho_\pm(x)$ and $q_\pm(y)$. In the second half, instead of the average quantities we plot sampled values (X, Y) from $Q_\pm(x, y)$ [depending on the value of $S(t)$]. This leads to qualitatively very similar output as with the full model in the left column of Fig. 11.

In the next two figures we show a quantitative comparison in Regime 3 between the full system and all three reduced models. In Fig. 12 we compare the stationary distribution of Y obtained by long-time simulations of the four models. We note an excellent agreement between the histogram from RM3 (solid green line) and that from the full model (solid black line). But this is not surprising since to get $q(y)$ with the RM3 we use the quasi-steady densities $q_\pm(y) = \int_{\Omega_\pm} Q_\pm(x, y)\, \mathrm{d}x$ (obtained in turn from a full model simulation with a reflecting condition at x_*), and the only output from RM3 is the proportion of time spend in the left/right wells. The histogram from RM2 (dot-dash blue line) captures the bimodality of Y in this

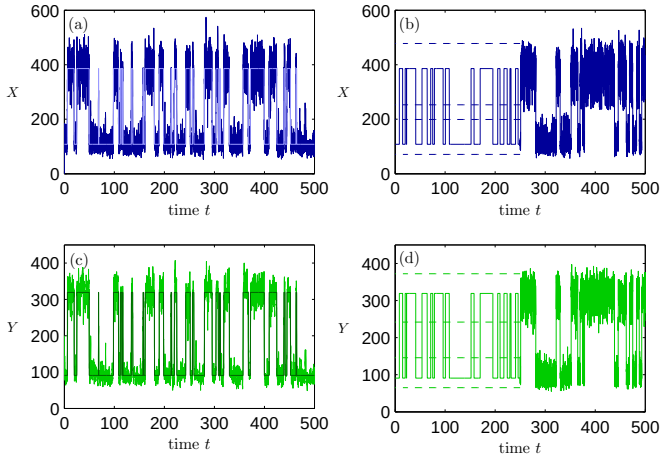

FIG. 11. Comparison of one realisation in Regime 3 of (a,c) the exact full system (6) and (b,d) the approximated reduced model 3 (19) using the NRM algorithm. We use $\tau_y = 0.25$, $X(0) = 0$ [*i.e.* $S(0) = S_-$ in (b,d)], $Y(0) = 200$, and the parameter values (7). In (a,c), the simulated quantities $X(t)$ and $Y(t)$ are shown in dark blue and light green, respectively. The two-valued curves in (a,c) are computed from (20). In (b,c), the approximate representation of $(X(t), Y(t))$ from the simulated $S(t)$ in shown as follows: in the first half, (20) and the 95% confidence intervals of $\rho_\pm(x)$ and $q_\pm(y)$ are plot as solid and dash lines, respectively. In the second half of (b) and (d), curves show samples of $\rho_\pm(x)$ and $q_\pm(y)$, respectively, taken every $\Delta t = 0.1$ from as per $S(t) = S_\pm$.

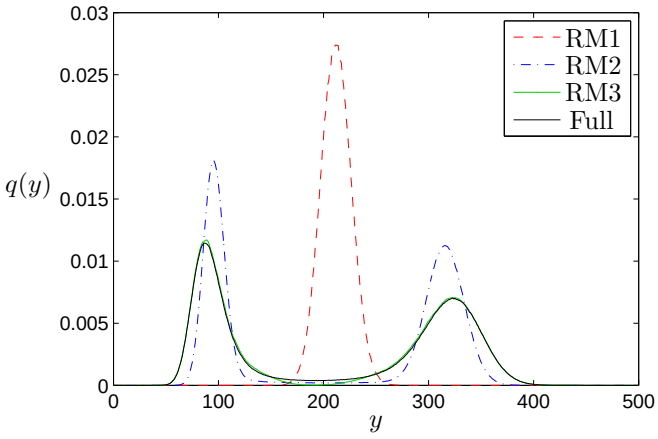

FIG. 12. Comparison of the marginal stationary density $q(y)$ in regime 3: full system (6) (solid black line) and approximate reduced models 1 (16) (dash red line), 2 (17) (dot-dash blue line), and 3 (19) (solid green line). All four histograms are computed by running the NRM algorithm up to $t = 10^6$ taking recordings every $\Delta t = 0.1$ (10^7 samples). We use $\tau_y = 0.25$ and the parameter values (7).

regime, but it gets the conditional means and variances substantially wrong. This is because X and Y are varying in the same timescale whereas RM2 supposes that X is much faster than Y. If the timescales in Regime 3 satisfied $\epsilon \ll \tau_y \ll \tau_s$, we would expect RM2 to give a good approximation to the full model dynamics.

Fig. 13 shows the mean $\mu_Y(t)$ and the standard deviation $\sigma_Y(t)$ of Y obtained from 10^5 realisations of the four models. We initialise the system in the left well ($S(0) = S_-$), using $X(0) \sim \rho_-(x)$ and $Y(0) \sim q_-(y)$ for the models that require explicit X or Y initialisation. The values of $\mu_Y(t)$ and $\sigma_Y(t)$ corresponding to RM3 can be calculated from $S(t)$. A simple calculation shows that they are given by

$$\begin{aligned} \mu_Y(t) &= \overline{Y}_- + \left(\overline{Y}_+ - \overline{Y}_-\right) s_+(t), \\ \sigma_Y^2(t) &= \sigma_{Y_-}^2 + \left[\sigma_{Y_+}^2 - \sigma_{Y_-}^2\right] s_+(t) \\ &\quad + \left[\overline{Y}_+ - \overline{Y}_-\right]^2 s_+(t)[1 - s_+(t)], \end{aligned} \tag{22}$$

where $s_+(t)$ is the estimated probability that $S(t) = S_+$ and $\sigma_{Y_\pm}^2 = \int (y - \overline{Y}_\pm)^2 q_\pm(y) \mathrm{d}y$. The mean μ_Y of both RM2 and RM3 captures well the mean of the full model, while RM1 misses the transient and jumps quickly to its stationary value (since it is taking X to be equilibrated instantly at $t = 0$). Similarly, the global standard deviation σ_Y of the full model is well approximated by RM2 and RM3, while RM1 underestimates (this is to be expected as we have eliminated the noise coming from the switches of X).

F. Estimation of mean-switching times

In reducing the fast variable X to a two-state Markov process in the reduced models 2 and 3 (corresponding to taking the limit $\delta \to 0$ as we will see in Appendix A 2), the key pieces of information we need to extract are the switching rates k_+ and k_-. These are the inverses of the mean transition (or escape) times T_+ and T_-. In this section we show how to accurately obtain these rates for the chemical system.

The estimation of the mean escape times for metastable processes is a classical problem that has received much attention in the literature.[31] As $\delta \to 0$ the escape times become exponentially small in δ and may be estimated by a variety of techniques in exponential asymptotics.[19,20] In some cases the process for X may be too complicated to estimate the mean escape times analytically, and a numerical estimate must be used.[32] However, when ϵ and δ are small but non-zero even the definition of escape becomes an issue: at what stage has the process X reached the other well?

For our simple bistable $X(t)$ process (8a), the mean time to reach any given point x_1, given we start at $X(0) = x_0$ is given exactly by[24]

$$\begin{aligned} T_-(x_0, x_1) &= \int_{x_0}^{x_1} \int_0^z \frac{\rho(s)}{d(z)\rho(z)} \mathrm{d}s\mathrm{d}z, \quad \text{with} \quad x_0 < x_1, \\ T_+(x_0, x_1) &= \int_{x_1}^{x_0} \int_z^\infty \frac{\rho(s)}{d(z)\rho(z)} \mathrm{d}s\mathrm{d}z, \quad \text{with} \quad x_0 > x_1. \end{aligned} \tag{23}$$

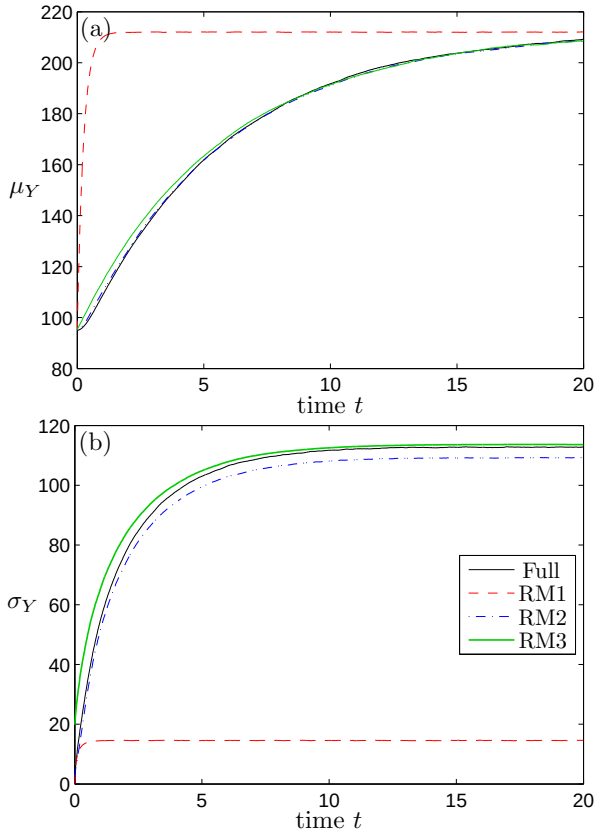

FIG. 13. Comparison of the time-dependent (a) mean $\mu_Y(t)$ and (b) standard deviation $\sigma_Y(t)$: full system (6) (solid black line) and approximate reduced models RM1 (16) (dash red line), RM2 (17) (dot-dash blue line), and RM3 (19) (solid green line). All curves are computed as the average over 10^5 realisations with initial conditions $X(0) \sim \rho_-(x)$ and $Y(0) = \overline{Y}_- \approx 95$. We use $\tau_y = 25$ and the parameter values in (7). The curves for RM3 are computed using (22).

The question is, what values do we choose for x_0 and x_1?

For x_0 we could choose the local maximum of ρ in the left-hand well (which we denote by x_-), or we could sample from the stationary distribution conditional on being in the left-hand well (which we denote by ρ_-). Since in the limit $\epsilon \ll \tau_s$ equilibration within a well is rapid by comparison to transitions between wells, we could in principle start with any value of x_0 in the left-hand well and we would obtain the same transition time to leading order.

For x_1 we could use x_*, the minimum of ρ, which satisfies

$$d'(x_*) - v(x_*) = 0.$$

We then need to double the mean first passage time to find the mean escape time, since a particle at x_* will return to the well it came from with probability one half. Alternatively, since equilibration within a well is rapid by comparison to transitions between wells, we may choose any x_1 which is sufficiently far from x_* so that immediate return to the left-hand well is unlikely. For example, we could choose the mean $(x_* + x_+)/2$, where x_+ is the local maximum of ρ in the right-hand well.[24]

In Table I we present values of $T_\pm$ [using (23)] and the corresponding results obtained by the SSA (with relative standard error of less than 1%, achieved with roughly 10^4 exits) for several choices of x_0 and x_1. We see that there is a small but significant variation in the switching times despite the fact that ϵ is quite small. To decide which of these times to use we consider here what properties we require of the reduced system (17a). If we denote $s_\pm(t)$ as the probability that $S(t) = S_\pm$ then $s_+(t)$ satisfies

$$\frac{\mathrm{d}s_+}{\mathrm{d}t} = k_- - (k_- + k_+)s_+, \tag{24}$$

where we have used the relation $s_- + s_+ = 1$ to eliminate s_-. If we initialise X in the left-hand well [by sampling from $\rho_-(x)$] then $s_+(0) = 0$, giving

$$s_+(t) = \frac{k_-}{k_- + k_+}\left(1 - e^{-(k_- + k_+)t}\right). \tag{25}$$

The stationary value $s_+ = \frac{k_-}{k_- + k_+} = 1 - \theta$, represents the proportion of time that X spends in the right-hand well on average. For S to be a good approximation to X this should be equal to the integral of $\rho(x)$ over the right-hand well. This gives us one relationship between k_+ and k_-, which determines the ratio between the mean switching times $T_\pm$. It can be shown that choosing k_+ and k_- which satisfy this constraint will ensure that the reduced process accurately captures the mean behaviour of Y.

To accurately capture the variance in Y we need to capture the rate of approach to this stationary solution accurately, i.e. we need to determine the time constant $\psi = 1/(k_- + k_+)$. We can avoid the difficulty of determining when a switch in X has occurred by considering the time dependent mean of both X and the reduced process S. If we can match the decay rates of these means, then we will have determined the rate constant ψ accurately. The mean of X under the reduced model is given by

$$\overline{X}(t) = \theta\overline{X}_- + (1-\theta)\overline{X}_+ - (1-\theta)(\overline{X}_+ - \overline{X}_-)e^{-t/\psi}, \tag{26}$$

where $\overline{X}_\pm = \int_{\Omega_\pm} x\rho_\pm(x)\,\mathrm{d}x$ are the mean values of X restricted to the right and left well respectively. The equilibrium value is $\overline{X}_\infty = \theta\overline{X}_- + (1-\theta)\overline{X}_+$. By comparing (26) to an ensemble of short-time stochastic simulations of the full process $X(t)$ we are able to get a good estimate ψ, which is then enough to determine k_- and k_+.

In Figure 14(a) we show the computed time evolution of $\overline{X}$, the mean of the full X process (obtained as the average of 10^5 runs). If from this we subtract the large time behaviour $\overline{X}_\infty$ and then take a logarithm we should obtain a straight line with gradient $-1/\psi$, as shown in Fig. 14(b). We fit a straight line to the part of the curve in which $\overline{X} - \overline{X}_\infty$ lies between 80% and 40% of its initial

TABLE I. Mean exit times T_- and T_+ to leave the left and right wells, respectively, as derived from numerical simulations and theory. We compare the results obtained using different start and end points.

T_- (left $\to$ right)				T_+ (right $\to$ left)			
x_0	x_1	Eq. (23)	SSA	x_0	x_1	Eq. (23)	SSA
$\rho_-(x_0)$	$(x_* + x_+)/2$	11.1190	11.1275	$\rho_+(x_0)$	$(x_- + x_*)/2$	11.5366	11.5573
x_-	$(x_* + x_+)/2$	11.4443	11.4610	x_+	$(x_- + x_*)/2$	11.8714	11.9482
0	$(x_* + x_+)/2$	11.6121	11.3253	1000[a]	$(x_- + x_*)/2$	11.9853	12.1354
0	x_*[b]	13.7263	13.7463	1000[a]	x_*[b]	11.2982	11.5974

[a] The starting point for the right to left transitions, which in theory should be $x = \infty$, is placed at $x = 1000$, where $\rho_+(1000) \sim 10^{-59}$. The relative error in the theoretical values using ∞ or 1000 are $O(10^{-4})$.

[b] The values $T_\pm$ using $x_1 = x_*$ are multiplied by two (both in the theory and SSA) since there is, roughly speaking, a 50% chance to fall either side once $X(t)$ reaches the unstable point x_*.

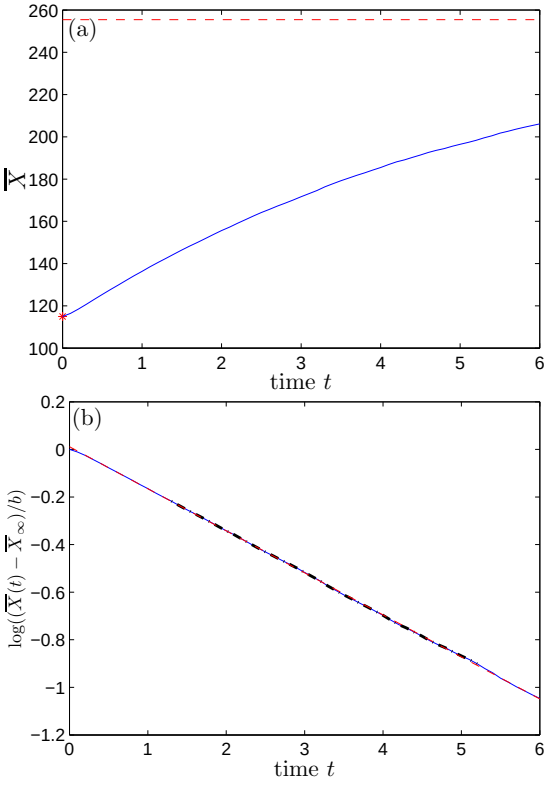

FIG. 14. Estimation of ψ from (26). (a) Time-evolution of the mean $\overline{X}(t)$ with $\overline{X}(0) = \overline{X}_-$ (solid blue line). The stationary value $\overline{X}_\infty$ is shown as a dash red line. The curve has been obtained from 10^5 SSA realisations of the full model (6). (b) Transformation $\log((\overline{X} - \overline{X}_\infty)/b)$ with $b = (\theta - 1)(\overline{X}_+ - \overline{X}_-)$ using $\overline{X}(t)$ in (a). The parameter ψ is estimated from the slope of this curve, using (26). The black dashed line shows the section we use for the linear fit (shown in as a dot-dashed red line).

value, in order to avoid any initial transients. The fit is good, and leads to the following values for θ, ψ and $T_\pm$:

$$\begin{aligned} \theta &= 0.4729, \qquad \psi = 5.6602, \\ T_- &= 10.7379, \quad T_+ = 11.9697. \end{aligned} \tag{27}$$

We see that these are in the same range as those in Table I. These are the values of $T_\pm$ that have been used in all simulations for RM2 and RM3 presented in Subsections III D and III E.

IV. CALCULATING EXTINCTION TIMES IN A PREDATOR–PREY SYSTEM

In this section we apply the model reduction methods developed above to an ecological model. In particular, we consider the probability of extinction of a population of predators when the prey undergoes a metastable stochastic process with bistability. This case study allows us to summarise the application of the method and, in particular, the two ingredients that one must extract from the original slow–fast system, namely the quasi-stationary densities and the mean switching times. Most importantly, this model invites us to push the method further by examining its performance with a system with absorbing states. We are thus interested in capturing the evolution of the system for low numbers. As a result, rather than using the continuous Fokker–Planck (FP) approximation for both species as we did for the chemical system, in this section we use the discrete description based on the backward master equation for the dynamics of the slow variable whose extinction we want to study.

We briefly give some background to our choice of model for this case study. One of the most fundamental questions in population biology concerns the persistence of species and populations, or conversely their risk of extinction.[33] Stochastic population models have become a common tool to investigate how the mean time to extinction depends on properties of the ecosystem. However, performing detailed mathematical investigations to understand how extinction risk depends on properties of the ecosystem is limited by the availability of tractable yet relevant models. A number of simple stochastic models of population dynamics have been used to study the effects of demographic processes on the mean time to extinction.[33] A small subset of these have looked at the situation in which a population exhibits bistability (*e.g.* Refs. 34–36). These are predominantly cases in which there simultaneously exists a positive and a zero abun-

TABLE II. Propensity functions.

Prey		Predator	
Transition	Propensity α_i	Transition	Propensity α_i
$X \to X+1$	$\alpha_1(x) = \lambda x - \lambda\frac{x^2}{\kappa}$	$Y \to Y+1$	$\alpha_3(y) = \tilde{\lambda}y$
$X \to X-1$	$\alpha_2(x) = \mu x + \frac{\beta z x^2}{1+\beta h x^2}$	$Y \to Y-1$	$\alpha_4(x,y) = \tilde{\mu}y + \frac{\tilde{\beta}y^2}{K(x)}$

dance attractor (locally stable steady states in the deterministic model). However, Palamara *et al.*[37] recently studied the mean time to extinction in simple predator–prey stochastic systems and reported a region of parameter space in which the prey population is bistable because there exists two alternative attractors with positive abundance, although they did not investigate that region for the mean time to extinction. We choose to use the model in Ref. 37 in the bistable regime as the basis for our case study because the population dynamics exhibited by the prey in this region are closely analogous to the bistable dynamics of X in the chemical reaction case (*e.g.* Fig. 15). To this we add a "predator" population whose abundance is influenced by the abundance of the prey, but does not itself influence the prey population (for example it survives on the direct by-products of the prey, such as dung beetles; while this is not strictly predation we will use the terms "predation" and "predator" here for clarity). In this case study we consider the probability of extinction of such a predator population when the prey undergoes a metastable stochastic process with bistability. We do not choose this example to represent any actual predator–prey system but rather to illustrate the advantages of using our model reduction techniques to enable the calculation of a property of domain-specific interest; in this case enabling efficient calculation of the probability of extinction of a population.

We identify three regimes as in the previous chemical example, depending on the relative timescales for prey-population switches and predator dynamics, and use our reduction method to measure the extinction rate of the predator population without having to simulate the computationally costly full predator–prey system. We consider the following birth and death model for the prey[37],

$$X \xrightarrow{k_1} 2X, \qquad\qquad X \xrightarrow{k_2} \emptyset, \tag{28a}$$

where

$$k_1(x) = \lambda\left(1 - \frac{x}{\kappa}\right), \qquad k_2(x) = \mu + \frac{\beta x z}{1+\beta h x^2}.$$

Here the reproduction of X corresponds to the Verhulst logistic model for population growth,[38] where λ is the intrinsic growth rate of the population and κ is referred to as the carrying capacity of the environment. This abstraction is commonly used to represent the limiting-effects of population density on population growth, through mechanisms such as resource limitation.[39] The population death rate includes a constant death rate μ (predator-free death rate) and a predation-induced death rate. We use a Holling's Type III functional response, representing a situation in which predators consume multiple prey items and switch in their feeding preference to predating on species X when it becomes particularly abundant.[39] The term βx is known as the attack rate, h is the prey handling time and z is the abundance of predators. Since we want this example to showcase the applicability of the method in a practical problem, here we do not introduce small parameters ϵ and δ as in the chemical system to make the characteristic timescales of the fast variable explicit.

We suppose that the predator pool for x (parameterised by z) is composed of many species, of which we pick one, denoted by y. We assume that changes in the predator y are negligible from the point of view of the prey, that is we can take z to be constant. We take the predator population to evolve according to the logistic model[40]

$$Y \xrightarrow{k_3} 2Y, \qquad\qquad Y \xrightarrow{k_4} \emptyset, \tag{28b}$$

where

$$k_3 = \tilde{\lambda}, \qquad k_4(x,y) = \tilde{\mu} + \frac{\tilde{\beta}y}{K(x)},$$

where $K(x)$ is a prey-dependent carrying capacity of the predator population. We will use $K(x) = \tilde{\kappa}(1 + x^\theta)$ in what follows. The second term in k_4 accounts for competition for resources: the death rate per individual increases with the predator population size and decreases with increasing availability of prey. Note therefore that, from an ecological standpoint, Y does not necessarily depend on X to maintain a positive population size, but its abundance is influenced by the availability of X. See table II for a summary of the propensity functions for each of the reactions.

In Fig. 15 we plot one run of the predator–prey system (28) for the following set of parameters

$$\begin{aligned} &\lambda = 1.5, && \mu = 0.5, && \beta = 0.015/z, && \kappa = 1000, \\ &h = 0.0055z, && \tilde{\lambda} = 5.5/\tau_y, && \tilde{\mu} = 3.5/\tau_y, && \tilde{\beta} = 5.5/\tau_y, \\ &\tilde{\kappa} = 0.002, && \theta = 2, \end{aligned} \tag{29}$$

for three different timescales τ_y for the predator. As before, we vary τ_y relative to the prey switching timescale τ_s [which we find to be $\tau_s = O(10^2)$ in the next subsection, see (31)] to move between regimes. Fig. 15(a) corresponds to $\tau_y = 10^4$ and is an example of a Regime

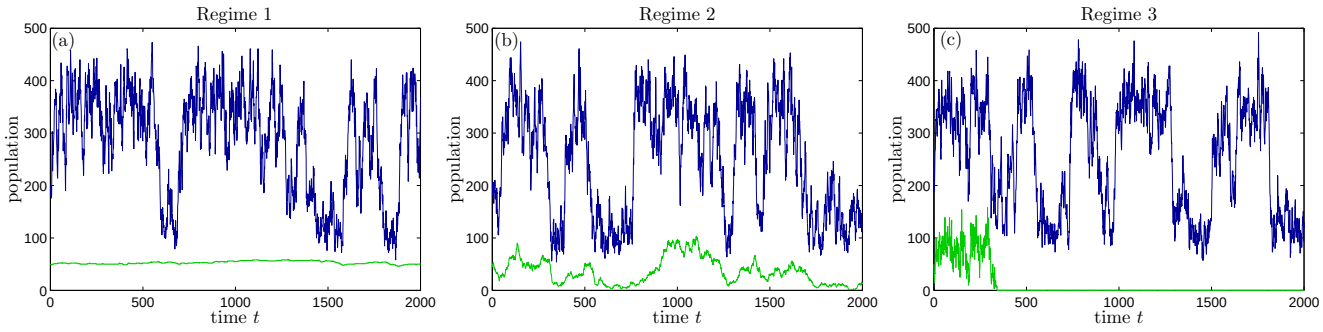

FIG. 15. Evolution of the prey $X(t)$ (blue lines) and the predator $Y(t)$ (green lines) from a single realisation of (28) obtained using the NRM algorithm. We use the parameter values (29), $X(0) = 200$, $Y(0) = 50$, and vary τ_y to change between regimes. (a) Regime 1: $\tau_y = 10^4$. (b) Regime 2: $\tau_y = 100$. (c) Regime 3: $\tau_y = 1$.

1 situation: the predator population hardly responds to fluctuations in the abundance of the prey. Fig. 15(b) shows a run with $\tau_y = 100$, which corresponds to Regime 2. Finally, in Fig. 15(c) the timescale of predator dynamics is set to be much shorter than the switching time of the prey by using $\tau_y = 1$, which fits with the definition of third regime.

We notice a new and interesting effect in this bistable system as we vary the timescale of the population of interest and change the regime: the predator population in Regime 3 [Fig. 15(c)] has become extinct before the final simulation time $T_f = 2000$. This is only one run of the stochastic system, but if we repeat the same experiment many times a clear pattern emerges: in Regime 3 the predator population reacts quickly to a low level of prey population and its chances of becoming extinct increase substantially. On the other hand, in Regime 1 the predator population only sees an average of the prey population and not its low and high levels and therefore the carrying capacity always stays at a level where time to extinction is long (since $\tilde{\lambda} > \tilde{\mu}$, giving a basic reproduction ratio above one.[38]) In what follows we use the slow–fast model reduction procedure to characterise the mean extinction time of the predator population as a function of its timescale relative to the switching time of the bistable prey population.

A. Quasi-stationary densities and mean switching times

Following the same procedure as for the chemical system, we compute the stationary density $\rho(x)$ of the prey (in this case it is a stationary distribution conditioned on the fact that extinction has not occurred, the so-called quasi-stationary density[38]).[41]

It is well-known that the Chemical Langevin approximation should not be used to predict the extinction rate or the quasi-stationary density near the extinction state, because it fails to correctly describe the very large fluctuations necessary to reach the absorbing state of zero particles.[42] However, here we focus on a parameter regime for which the metastable prey population is large and the relaxation time to its quasi-steady density $\rho(x)$ is extremely small compared to its mean time to extinction. We thus find that the stationary solution of the Fokker–Planck equation with a reflecting boundary condition at $x = 0$, equation (13), gives an accurate estimate of the prey quasi-stationary density. Methods to determine this density more accurately in the region near extinction are available (see *e.g.* Ref. 43). The drift and diffusion coefficients are

$$v(x) = x[k_1(x) - k_2(x)], \quad d(x) = \frac{1}{2}x[k_1(x) + k_2(x)], \tag{30}$$

The resulting stationary density, which we denote again by $\rho(x)$, is shown in Fig. 16(a). We observe that $\rho(x)$ is bimodal, indicating that the prey population is metastable and will switch between its two favourable states $x_- = 108.5$ and $x_+ = 352.1$.

In Fig. 16(b) we show the histograms of the quasi-stationary predator distributions $q_\pm(Y)$ conditioned on the event of low and high prey levels respectively (X lower or higher than $x_* = 205.0$), for the parameter values in (29) and $\tau_y = 100$. We use the capital letter Y to emphasise that we are not taking the continuum limit here (assuming Y is large is not appropriate if we are interested in extinction)

Next, we evaluate the mean-switching times between the low and high prey-population levels as in section III F. The mean time X spends in each well is estimated from the stationary density shown in Fig. 16(a). The evolution of the mean of X is calculated by averaging over 10^6 realisations up to $T_f = 80$. As before, we fit an exponential decay to the portion of the curve in which $\overline{x} - \overline{x}_\infty$ lies between 80% and 40% of its initial value. From this analysis we obtain the following values for θ, ψ and $T_\pm$:

$$\begin{aligned} \theta &= 0.3004, & \psi &= 76.8366, \\ T_- &= 109.8315, & T_+ &= 255.7689. \end{aligned} \tag{31}$$

We note that the mean switching times $T_\pm$ set the size of the switching timescale τ_s. Having obtained the quasi-

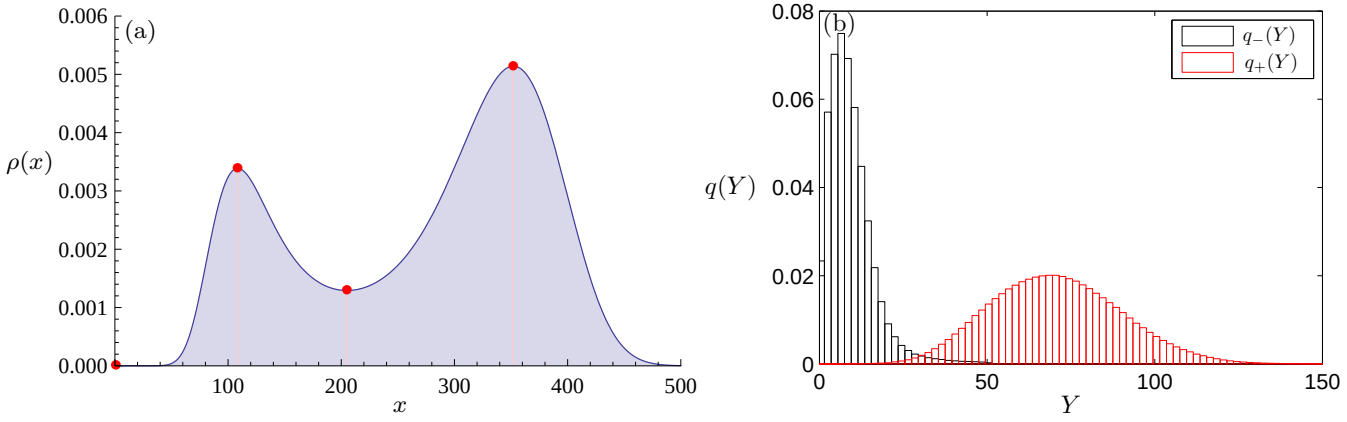

FIG. 16. (a) Quasi-stationary marginal density $\rho(x)$ computed from (13) using the drift and diffusion coefficients in (30). (b) Quasi-stationary marginal densities $q_\pm(Y)$ of Y from SSA of (28) conditional on the prey being at low ($-$, black histogram) and high ($+$, red histogram) levels respectively, obtained from 10^7 steps of the NRM algorithm. We use $\tau_y = 100$ and the parameter values (29).

steady densities $\rho_\pm(x)$ of the fast metastable species and its mean switching times $T_\pm$, we are ready to apply the reduced model approximations to simulate the predator–prey system. Our goal is to estimate the mean time to extinction (MTE) of the predator population.

B. Mean times to extinction

Since we are interested in extinction we do not use a continuum approximation for Y but retain a discrete approximation. We denote by $\mathcal{T}_n$ the MTE of Y given that $Y(0) = n$. It can be determined exactly via the backward master equation.[15] We will use this approach on the original model as well as on the reduced ones, as detailed below.

1. Full system

We denote by $\mathcal{T}_n^m$ the two-dimensional MTE of the full system (28) given that $X(0) = m$ and $Y(0) = n$. It can be easily shown [see analogous one-dimensional example in (35) below] that it satisfies

$$\begin{aligned} -1 = {} & \alpha_1(m)\mathcal{T}_n^{m+1} + \alpha_2(m)\mathcal{T}_n^{m-1} \\ & + \alpha_3(n)\mathcal{T}_{n+1}^m + \alpha_4(m,n)\mathcal{T}_{n-1}^m \\ & - [\alpha_1(m) + \alpha_2(n) + \alpha_3(n) + \alpha_4(m,n)]\mathcal{T}_n^m, \end{aligned} \tag{32}$$

with absorbing left-boundary conditions $\mathcal{T}_0^m = 0$, $\mathcal{T}_n^0 = t_n$ and a boundary condition at $m = \kappa$, $\alpha_2(n)(\mathcal{T}_n^{m-1} - \mathcal{T}_n^m) = 1$, using that $\alpha_1(\kappa) = 0$. Here t_n is the mean time to extinction when there is no prey, which can be found by solving Eq. (35) below replacing $\overline{\alpha}_4(n)$ by $\alpha_4(0,n)$. In principle this defines an infinite set of difference equations for $\mathcal{T}_n^m$ for $n \geq 1$ and $0 \leq m \leq \kappa$. However, noting that $\overline{\alpha}_4 \gg \alpha_3$ for large n (due to the quadratic term in $\overline{\alpha}_4$, representing the competition), we can introduce an artificial right boundary condition at $n = N$ for large N and use a similar argument to that of $m = \kappa$ to give $\alpha_4(m,N)(\mathcal{T}_{N-1}^m - \mathcal{T}_N^m) = -1$. This boundary condition can be imposed by adopting the convention that $\alpha_3(N) = 0$. Finally, to obtain $\mathcal{T}_n$ we integrate $\mathcal{T}_n^m$ against the quasi-stationary density $\rho(x)$ in Fig. 16(a).

2. Reduced model 1 ($\tau_s \ll \tau_y$)

The appropriate reduced model for the system in Regime 1, analogous to RM1 in (16), is

$$\begin{aligned} & Y \to Y+1: \ \alpha_3(Y) = \tilde{\lambda} Y, \\ & Y \to Y-1: \ \overline{\alpha}_4(Y) = \tilde{\mu} Y + \frac{\tilde{\beta} Y^2}{\overline{K}}, \quad \frac{1}{\overline{K}} = \int_0^\infty \frac{\rho(x)}{K(x)}\,\mathrm{d}x. \end{aligned} \tag{33}$$

To obtain the MTE from (33), we consider what can the predator population do in the first short time interval δt:

$$\begin{aligned} \mathcal{T}_n - \delta t = {} & \alpha_3(n)\delta t \mathcal{T}_{n+1} + \overline{\alpha}_4(n)\delta t \mathcal{T}_{n-1} \\ & + [1 - \alpha_3(n)\delta t - \overline{\alpha}_4(n)\delta t]\,\mathcal{T}_n. \end{aligned} \tag{34}$$

Hence we obtain

$$\alpha_3(n)\mathcal{T}_{n+1} - [\alpha_3(n) + \overline{\alpha}_4(n)]\mathcal{T}_n + \overline{\alpha}_4(n)\mathcal{T}_{n-1} = -1. \tag{35}$$

Similarly as before, we solve (35) together with $\mathcal{T}_0 = 0$ and $\alpha_3(N) = 0$ for N large.

3. Reduced model 2 ($\tau_s \sim \tau_y$)

Analogously to RM2 defined in (17a), the reduced model in Regime 2 has a death rate dependent on

whether the prey is at a low or high population level:

$$Y\to Y+1\colon\ \alpha_3(Y)=\tilde{\lambda}Y,$$
$$Y\to Y-1\colon\ \alpha_4^{\pm}(Y)=\tilde{\mu}Y+\frac{\tilde{\beta}Y^2}{\overline{K}_{\pm}},\quad \frac{1}{\overline{K}_{\pm}}=\int_{\Omega_{\pm}}\frac{\rho_{\pm}(x)}{K(x)}\,\mathrm{d}x, \tag{36a}$$

where $\rho_{\pm}(x)$ are the quasi-stationary prey densities taken from $\rho(x)$ conditioned on $X\in\Omega_{\pm}$ (left or right wells). The switches between low- and high-level prey population obey the reactions

$$S_-\underset{k_+}{\overset{k_-}{\rightleftharpoons}}S_+,\qquad k_{\pm}=1/T_{\pm} \tag{36b}$$

with $T_{\pm}$ given in (31). This time the MTE must take into account the initial state of the boolean variable $S(t)$. Denote $\mathcal{T}_n^{\pm}$ the MTE given that $S(0)=S_{\pm}$ and $Y(0)=n$, respectively. Then, following a similar argument as before, $\mathcal{T}_n^{\pm}$ obey

$$\begin{aligned}&\alpha_3(n)\mathcal{T}_{n+1}^{\pm}+\alpha_4^{\pm}(n)\mathcal{T}_{n-1}^{\pm}+k_{\pm}\mathcal{T}_n^{\mp}\\&\quad-[\alpha_3(n)+\alpha_4^{\pm}(n)+k_{\pm}]\mathcal{T}_n^{\pm}=-1,\end{aligned} \tag{37}$$

with $\mathcal{T}_0^{\pm}=0$ and by convention $\alpha_3(N)=0$. Again this results in a closed set of equations, now with $2N$ equations and unknowns.

4. Reduced model 3 ($\tau_y\ll\tau_s$)

Recall that the reduced model 3 does not keep track of explicit Y dynamics but instead only those of the switching variable S. In order to extract a MTE in this reduced model, we must introduce a new pair of reactions to represent Y-extinction from each of the wells:

$$S_-\underset{k_+}{\overset{k_-}{\rightleftharpoons}}S_+,\qquad S_-\xrightarrow{r_-}\emptyset,\qquad S_+\xrightarrow{r_+}\emptyset, \tag{38}$$

where $k_{\pm}=1/T_{\pm}$ as before and $\emptyset$ means "Y is extinct". Here the rates of predator extinction from the right and left wells are $r_{\pm}=1/\Pi_{\pm}$, where $\Pi_{\pm}$ are the MTE of Y starting from its quasi-stationary distribution $q_{\pm}(Y)$ and conditional of X staying in a given well for all times. To evaluate $\Pi_{\pm}$, we solve the equation for the MTE in the full system (32) but with a reflective boundary condition at $X=x_*$, and integrate the resulting matrix $\mathcal{T}_n^m$ against the quasi-stationary densities $Q_{\pm}(m,n)$ (computed similarly as for the chemical system). For any τ_y one finds that $\Pi_-\ll\Pi_+$ since it is much more likely for the predator to become extinct when the prey population is in the left well than when it is in the right well [as can be seen in Fig. 16(b)]. From (38) it is easy to show that the MTE $\mathcal{T}^{\pm}$ of the reduced model in Regime 3 given that $S(0)=S_{\pm}$ is

$$\mathcal{T}^{\pm}=\frac{k_-+k_++r_{\mp}}{k_+r_-+k_-r_++r_-r_+}. \tag{39}$$

The approximation given by RM3 of the MTE $\mathcal{T}_n$ if we don't know where $X(t)$ started [$X(0)\sim\rho(x)$] is given by $\mathcal{T}=\theta\mathcal{T}^-+(1-\theta)\mathcal{T}^+$. Note that the result in RM3 does not depend on the initial value of Y.

5. Results

Using the exact expressions defined above for the MTE of $Y(t)$ in the original system as well as in each reduced model, we can look at the accuracy of each reduced model relative to the exact full model as a function of τ_y. We show the results in Fig. 17 for $Y(0)=30$ and the parameter values (29) and (31). We choose the right boundary at $n=N$ such that the results are insensitive to N; we find that $N=150$ is a good choice. The region of validity of each reduced model is set by the switching scale $\tau_s\sim10^2$ (*e.g.*, when $\tau_y\sim10^2$ the system is in Regime 2).

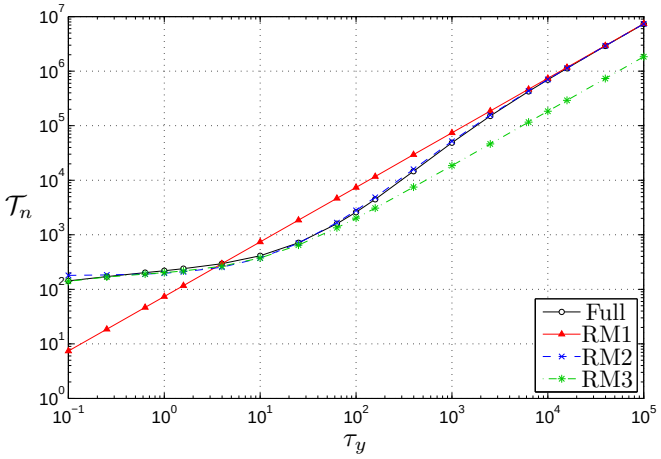

FIG. 17. Theoretical mean time to extinction $\mathcal{T}_n$ of Y as a function of τ_y. Curves obtained from the reduced models RM1 (35), RM2 (37) and RM3 (39), and the exact full system MTE (32). We use $n=30$, $N=150$ and the parameter values (29) and (31). The system (28) is in Regime 2 for τ_y around 100, and moves towards Regime 1 (3) for larger (smaller) values of τ_y.

We find that the MTE computed from the full model increases with τ_y roughly linearly (as expected) from $\tau_y\sim O(10^2)$, but that for faster timescales it becomes less sensitive to τ_y (see black line with circles in Fig. 17). This is because, when the system is in Regime 3 and $Y(t)$ is fast, the main contribution to the $\mathcal{T}$ is the time it takes for a switch in X from the right to the left well to occur (*i.e.* $T_+=1/k_+$ on average) multiplied by the probability that X started in the right well, which is $1-\theta$. Once the switch to the low-level prey has occurred, the extinction of Y is almost instantaneous for $\tau_y\to0$. This is because this extinction occurs at a rate r_- [see (38)] which scales like $1/\tau_y$. If we do this simple calculation, we find that $\mathcal{T}_n$ should tend to $T_+(1-\theta)=178.94$ as

TABLE III. Mean time to extinction $\mathcal{T}_n$ of Y with $Y(0) = n = 30$ obtained from theory and simulations of the full system and the reduced models in three parameter τ_y regimes. We use the same parameters and equations as in Fig. 17. For the simulated MTE, we run iterations of the SSA until the standard error in the estimate of $\mathcal{T}_n$ is below 1% and indicate the execution times in parenthesis.

Model	Regime 1 ($\tau_y = 10^4$)		Regime 2 ($\tau_y = 100$)		Regime 3 ($\tau_y = 1$)	
	Th.	SSA	Th.	SSA	Th.	SSA
Full	6.968×10^5	6.545×10^5 ($\approx$97h)[a]	2600.5	2566.7 (1097s)	219.24	219.58 (464s)
RM1	7.394×10^5	7.313×10^5 (11s)	7394.5	7333.2 (11s)	73.945	73.975 (10s)
RM2	7.043×10^5	7.058×10^5 (13s)	2790.8	2796.9 (8s)	197.64	199.98 (162s)
RM3	1.840×10^5	1.825×10^5 (0.9s)	2015.3	2018.1 (0.01s)	202.25	203.709 (0.005s)

[a] Value of $\mathcal{T}_{30}$ for $\tau_y = 10^4$ computed from 400 runs of (28) (relative error of 6%). CPU time estimated from the execution time of 400 runs and the 10^4 runs required to obtain values with 1% relative error (estimated from the number of rounds required for $\tau_y = 10^4$ in the reduced models).

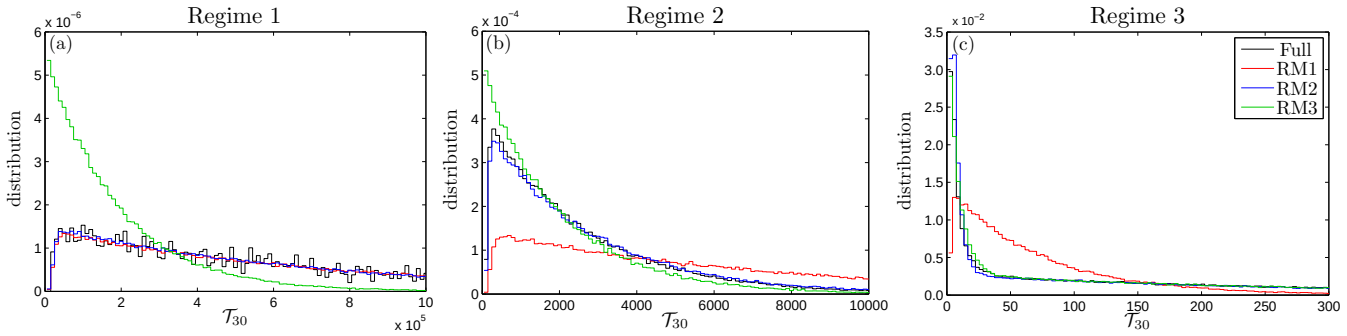

FIG. 18. Distribution of the time to extinction $\mathcal{T}_n$ with $Y(0) = n = 30$ for the parameters in (29) and $\tau_y = 10^4, 100$ and 1 (corresponding roughly to a system in Regime 1, 2 and 3 respectively). Histograms computed by 10^5 rounds of the SSA of the full system (28), and reduced models 1 (33), 2 (36) and 3 (38), respectively [except for the full system in (a), in which only 3300 rounds are used].

$\tau_y \to 0$, using (31). This value is consistent with the results plotted in Fig. 17. We see that the approximation of the MTE given by RM2 agrees very well with the exact MTE from the full system throughout the whole range of timescales τ_y except for very small scales where a small error can be perceived. As expected, the approximation to the MTE of RM1 is good for slow timescales τ_y such that the system is in Regime 1 ($\tau_y \to \infty$), but very poor for fast timescales in Regime 3 when the MTE is highly dependent on the bistable prey behaviour. Conversely, the approximation of RM3 is good in Regime 3 (τ_y small), but underestimates the MTE of Y when the predator is very slow (and is thus not as sensible to switches to the low-level in the prey population).

In Table III we compare the theoretical results of Fig. 17 with those obtained from multiple repetitions of the SSA. We do this to have an idea of the computational time reduction that each reduced model gives relative to the original system. We choose three values of the timescale τ_y such that they are representative of each regime: $\tau_y = 10^4$ for Regime 1, $\tau_y = 100$ for Regime 2, and $\tau_y = 1$ for Regime 3. To compute $\mathcal{T}_n$ numerically, we run the SSA of the full system (28) and the three reduced models RM1 (33), RM2 (36), and RM3 (38). We initialise the system with $Y(0) = 30$ and $X(0) \sim \rho(x)$ [equivalently, for RM2 and RM3, $S(0) = S_-$ with probability θ] and run the simulation until $Y(t) = 0$. We record the final time and repeat the process until the standard error in $\mathcal{T}_n$ is below 1%. The twelve simulation values of $\mathcal{T}_{30}$ in Table III required $O(10^4)$ runs to achieve such relative error.

Finally, in Fig. 18 we show the histograms of the time to extinction $\mathcal{T}_{30}$ obtained from the simulations. We compare the distribution of $\mathcal{T}_{30}$ of the full system with those from the reduced models. We see that the Regime 2 model does a good job of capturing the distribution for the whole range of values of τ_y. In contrast the reduced models 1 and 3 fail to capture the distribution of $\mathcal{T}_{30}$ near the origin when used outside of their regime of validity. This would be important if, for instance, we were interested in the probability that the predator population became extinct within one year.

V. SUMMARY AND DISCUSSION

In this paper we developed a model reduction technique for slow–fast stochastic systems with metastability. This is a generalisation of previous approaches based on stochastic averaging principles,[17] when the average of the

fast process switches between different quasi-stationary densities instead of relaxing to a unique stationary measure. Interestingly, we found that we can still accurately describe the behaviour of metastable slow–fast systems while improving numerical efficiency substantially by introducing a switch variable to describe the metastable process and using stochastic averaging separately in each of the basins of attraction.

We focused on a simple class of metastable slow–fast stochastic systems consisting of two species, one of which (X) is fast with exponentially long bistable switches. The second species (Y), which we assume is the species of interest, has dynamics that are coupled to the bistable species. We identified three dynamical regimes that led to different model reductions in both of our two-species metastable systems. When the dynamics of Y are much slower than the switches in X (and thus not affected by the bistable behaviour) then the reduced model RM1 is equivalent to the standard fast-variables reduction used in monostable systems.[17] However, in the other two regimes, when the evolution of Y is of the same order or faster than the switches in X, then the switches in X must be taken into account. Our key insight is that two ingredients are required to build a reduced model for slow–fast system with metastability in general. The first ingredient is knowledge of the transition rates between the different basins of attraction for the fast variables (which could of course be more than two in other situations). The second ingredient is the need for estimates of the quasi-stationary densities of the fast variables, conditioned on being in each basin of attraction. This is the procedure used by Markov state models for the modelling of molecular kinetics, with powerful yet computationally expensive tools available to do this in high-dimensional systems, more suited for parallel computation.[44]

In the chemical example, we showed how asymptotic approximations of the transition rates can be obtained from the Fokker–Planck equation, as well as from a short SSA run of the full system. In general applications the latter approach is likely to be more feasible because the Fokker–Planck equations can only be used in a limited set of analytically tractable models. However, as seen in section III F, care must be taken in defining what it means for the system to have switched to another favourable state. Switching rates have been inferred from other model parameters in a related way for ion channels.[45,46] For both the chemical and the ecological examples, we obtained the quasi-stationary densities of the fast variables from both the Fokker–Planck equation equation and the SSA results: again, the limitations of using the Fokker–Planck equation for more general purposes apply. However, once the conditional quasi-steady densities and the switching rates have been obtained, then we can compute the effective rates for the slow variables conditioned on each basin and the transition fluxes between sub-models.

In this work we assumed a known slow–fast model with fixed parameters. One natural extension to our study is to systems in which the parameters affecting the regime may change through time. For example, if τ_y changed through time, then one may be inclined to dynamically change the reduced model being employed. This would require building a set of rules or boundaries to discern between regimes (precomputed), and checking during a simulation whether any of these have been crossed. However, the extra computational effort involved in such a process is likely to be more costly than simply using the reduced model 2 (RM2) throughout the simulation. This model interpolates between Regimes 1 and 3 and is valid in the whole parameter regime for τ_y (*e.g.* see Fig. 17). To our knowledge this is the first method that can be used in the whole spectrum of τ_y in the class of systems considered here, thus extending the method in Ref. 28 to cases with more than one stable state, while still keeping it simple (a one-dimensional model plus the switch variable), in contrast to more complicated fully two-dimensional models appropriate in Regime 3.[47] A potentially more interesting situation arises if the (quasi-)stationary density of the fast variable $\rho(x)$ changes dynamically; maybe even through a two-way coupling with y. In this case, the RM2 as presented would need to be modified to since one needs to update dynamically the mean switching times $T_\pm$ and the conditional densities $\rho_\pm(x)$. Moreover, it could happen that $\rho(x)$ evolved from a bimodal to unimodal shape, in which case the system would move into into the monostable regime. A possible strategy would be to establish a timestep Δ_y such that $\epsilon \ll \Delta_y \ll \tau_y$, and every Δ_y stop the simulation, do a short run of the full dynamics to re-evaluate $T_\pm^{\hat{y}}$ and $\rho_\pm^{\hat{y}}(x)$, parameterised by current value of $Y = \hat{y}$.

We applied our techniques to a metastable system with an absorbing state; the predator–prey stochastic model with bistable prey (X), and showed that the reduced models can be used to predict the mean times to extinction (MTE) of the predator population Y in an accurate and efficient way. Fundamentally, this showed that the method can be applied more generally to two-dimensional systems, even with absorbing states. This implies that the specific forms of the rates and the coupling between X and Y do not matter for the application of our model reduction techniques. Moreover, the model reduction has provided some new insights into how the routes to extinction of the predator change as we varied its relative timescale to the prey switching and its quasi-stationary density changed from unimodal to bimodal. In particular, the reduced model RM2 can identify the huge difference between the MTE of Y when the prey population X is at its low or high level, and, most importantly, how this fact together with the bistable process translates into the overall MTE. The specific application of our approximations to enable the efficient calculation of the mean time to extinction in Regime 2 is, to our knowledge, new to the literature and clearly illustrates the value of the approximations of reduced model 2 (explicit Y with a switch variable for X) over the original system to estimate the mean time to extinction. It was fortunate for

these equations that we could assess our estimates with the backward master equation; however, this will rarely be possible for higher-dimensional ecological models. It is in such situations that having insight into the appropriate reduction method is most valuable, giving one the ability to select the appropriate stochastic reduced model to obtain accurate approximations efficiently. In high dimensional systems with several variables displaying intrinsic metastable behaviour, one could use a combination of our method presented here for the MTE and the computational approaches discussed in Ref. 44 to split the space into basins of attraction and approximate the transitions times.

Finally, this work could be extended to enable long-term predictions of complex real bistable systems. This could, for example, enable new insights into our understanding of critical transitions and our ability to predict them.[9] In recent years there has been a lot of effort in investigating time-series of systems prone to critical transitions in an attempt to enable predictions of transitions.[48] Could we detect the dynamical regime of a bistable system from its dynamics without knowing the parameters or model? If we could then it is possible to imagine using similar ingredients to those used here to infer the mean switching times characterising the metastable processes and the reaction time (after a switch) for the variables of interest.

ACKNOWLEDGMENTS

We thank J. M. Newby, G. M. Palamara, and Y. G. Kevrekidis for useful discussions on aspects of this article. We also thank M. Geissbuehler and T. Lasser for their *Morgenstemning* colormap.[49] M.B. is partially funded by the EPSRC (EP/I017909/1) and Microsoft Research, Cambridge and by St John's College, Oxford, in the form of a Junior Research Fellowship.

Appendix A: Perturbation analysis of the three bistable regimes

In this appendix we provide formal derivations of the three reduced models in the main text using a perturbation analysis at the level of the Fokker–Planck equation for the joint probability density $P(x, y, t)$. In A 1 we consider the asymptotic regime $\tau_s \ll \tau_y$ (Regime 1), resulting in the reduced model 1. The more interesting reduced model 2, appropriate when $\tau_s \sim \tau_y$, is derived in A 2 using a WKB perturbation method. Finally, in A 3 we discuss the regime when $\tau_y \ll \tau_s$ and the reduced model 3 is appropriate.

We consider the Fokker–Planck equation (10) for the joint probability density $P(x, y, t)$. The main timescale of interest is that of the evolution of Y, so let us rescale time with τ_y to give

$$\frac{\partial P}{\partial t}(x,y,t) = \frac{\tau_y}{\epsilon}\frac{\partial}{\partial x}\left\{\frac{\partial}{\partial x}[d(x)P] - v(x)P\right\} + \frac{\partial}{\partial y}\left\{\frac{\partial}{\partial y}[D(x,y)P] - V(x,y)P\right\}, \tag{A1}$$

where $\tau_y/\epsilon \gg 1$.

1. Regime 1: $\epsilon \ll \tau_s \ll \tau_y$

First we consider equation (A1) for short times such that $t = O(\epsilon)$. We define the fast time $\tilde{t}$ as $t = (\epsilon/\tau_y)\tilde{t}$ and write $\tilde{P}(x, y, \tilde{t}) = P(x, y, t)$, to give, at leading order,

$$\frac{\partial \tilde{P}^{(0)}}{\partial \tilde{t}} = \frac{\partial}{\partial x}\left\{\frac{\partial}{\partial x}[d(x)\tilde{P}^{(0)}] - v(x)\tilde{P}^{(0)}\right\}. \tag{A2}$$

As $\tilde{t} \to \infty$, the solution of (A2) converges to

$$\tilde{P}^{(0)}(x,y) = \rho(x)q(y), \tag{A3}$$

where $\rho(x)$ is the normalised steady solution of (A2) [given by (13)], and

$$q(y) = \int_0^\infty P(x,y,0)\,\mathrm{d}x.$$

Now we move back to $O(1)$ times and consider equation (A1). Expanding $P \sim P^{(0)} + \epsilon/\tau_y\, P^{(1)} + \cdots$, gives, at leading-order,

$$\frac{\partial}{\partial x}\left\{\frac{\partial}{\partial x}[d(x)P^{(0)}] - v(x)P^{(0)}\right\} = 0. \tag{A4}$$

Thus

$$P^{(0)}(x,y,t) = C(y,t)\rho(x), \tag{A5}$$

where $C(y, t)$ is arbitrary at this stage. Matching this solution for long times with the short times solution (A3) gives $C(y, 0) = q(y)$. At the next order equation (A1) gives

$$\frac{\partial P^{(0)}}{\partial t} = \frac{\partial}{\partial x}\left\{\frac{\partial}{\partial x}[d(x)P^{(1)}] - v(x)P^{(1)}\right\} + \frac{\partial}{\partial y}\left\{\frac{\partial}{\partial y}[D(x,y)P^{(0)}] - V(x,y)P^{(0)}\right\}. \tag{A6}$$

Integrating this equation with respect to x (using no-flux boundary conditions at $x = 0,\ \infty$) gives the following solvability condition for C:

$$\frac{\partial C}{\partial t}(y,t) = \frac{\partial}{\partial y}\left\{\frac{\partial}{\partial y}\left[\overline{D}(y)C(y,t)\right] - \overline{V}(y)C(y,t)\right\} \tag{A7}$$

where

$$\overline{F}(y) = \int F(x,y)\rho(x)\,\mathrm{d}x, \qquad \text{for } F = D,\ V.$$

Finally, since

$$C(y,t) = \int_0^\infty P(x,y,t)\,\mathrm{d}x$$

we see that $C(y,t)$ is the marginal density for $Y(t)$, so that (A7) gives the evolution of the probability density function of a reduced process for $Y(t)$ where the fast variable $X(t)$ has been averaged out. Going back to the original time variable we can then write

$$\frac{\partial p}{\partial t}(y,t) = \frac{1}{\tau_y}\frac{\partial}{\partial y}\left\{\frac{\partial}{\partial y}\left[\overline{D}(y)p\right] - \overline{V}(y)\right\} \tag{A8}$$

Equation (A8) is the Fokker–Planck equation associated with the reduced stochastic model (16) appropriate for Regime 1 when Y is approximated by a continuous random variable. This asymptotic reduction relies on the fact that the stochastic process for X reaches steady state on a timescale which is faster than the timescale for the evolution of Y.

2. Regime 2: $\epsilon \ll \tau_s \sim \tau_y$

To analyse Regime 2 asymptotically we need to ensure that X is metastable, and quantify the switching time. We suppose then that $1/\delta$, the typical equilibrium value of X, is large, and scale the rate constants as in (7). We set $\hat{x} = \delta x$, where $\hat{x}$ is $O(1)$ as $\delta \to 0$. We will see that $\tau_s \sim e^{r/\delta}$ for some constant $r > 0$. In terms of the new scaled variables equation (A1) becomes

$$\begin{aligned}\frac{\partial P}{\partial t} = &\frac{\tau_y}{\epsilon}\frac{\partial}{\partial \hat{x}}\left\{\frac{\partial}{\partial \hat{x}}[\delta\hat{d}(\hat{x})P] - \hat{v}(\hat{x})P\right\} \\ &+ \frac{\partial}{\partial y}\left\{\frac{\partial}{\partial y}[D(\hat{x},y)P] - V(\hat{x},y)P\right\},\end{aligned} \tag{A9}$$

where $\hat{d}(\hat{x}) = \delta d(x)$, etc. With this scaling we see that the diffusion $\delta\hat{d}(\hat{x})$ of X is weaker than the drift $\hat{v}(\hat{x})$ when $\delta \ll 1$, which is the reason that switches do not occur frequently ($\tau_s \gg \epsilon$); the parameter δ makes explicit this separation of timescales between drift of X and switches in X (this is the weak noise limit[50]).

It is convenient to write (A9) in the form

$$\hat{\epsilon}\frac{\partial P}{\partial t} + \hat{\epsilon}\frac{\partial J_y}{\partial y} = \mathcal{L}_\delta, \tag{A10}$$

where $\hat{\epsilon} = \epsilon/\tau_y$, $J_y = -\frac{\partial}{\partial y}[D(\hat{x},y)P] + V(\hat{x},y)P$ and $\mathcal{L}_\delta$ is a linear operator acting on the variable $\hat{x}$ only:

$$\mathcal{L}_\delta P \equiv \frac{\partial}{\partial \hat{x}}\left(\frac{\partial}{\partial \hat{x}}[\delta\hat{d}(\hat{x})P] - \hat{v}(\hat{x})P\right). \tag{A11}$$

The marginal stationary density in $\hat{x}$ [$\rho(x)$ in (13)] corresponds to the zero-eigenvalue eigenfunction $\mathcal{L}_\delta$, that is $\mathcal{L}_\delta\rho(\hat{x}) = 0$, giving, in the scaled variables,

$$\rho(\hat{x}) = \frac{A}{\hat{d}(\hat{x})}\exp\left[\frac{1}{\delta}\int_0^{\hat{x}}\frac{\hat{v}(s)}{\hat{d}(s)}\,\mathrm{d}s\right], \tag{A12}$$

where A is the normalisation constant.

The novelty in this bistable problem is the coupling of X with Y, that is, in the term J_y in (A10). We sketch the following calculation since it follows closely that of Refs. 19 and 20. By using a WKB approximation for small δ (with $\hat{\epsilon} \ll \delta$) we find solutions to (A10) of the form

$$P \sim \begin{cases} A_-\phi(\hat{x}) & \hat{x} < \hat{x}^v_*, \\ \frac{1}{2}\phi(\hat{x}^v_*)e^{-\gamma_*(\hat{x}-\hat{x}^v_*)^2/2\delta}\left[A_+ + A_- + (A_+ - A_-)\mathrm{erf}\left(\sqrt{\frac{|\gamma_*|}{2\delta}}(\hat{x}-\hat{x}^v_*)\right)\right] & \hat{x} \approx \hat{x}^v_*, \\ A_+\phi(\hat{x}) & \hat{x} > \hat{x}^v_*. \end{cases} \tag{A13}$$

where

$$\phi(\hat{x}) = \frac{e^{u(\hat{x})/\delta}}{\hat{d}(\hat{x})}, \quad u(\hat{x}) = \int_0^{\hat{x}}\frac{\hat{v}(s)}{\hat{d}(s)}\mathrm{d}s, \quad \gamma_* = -\frac{\hat{v}'(\hat{x}^v_*)}{\hat{d}(\hat{x}^v_*)}, \tag{A14}$$

and $\hat{x}^v_*$ is the turning point at which $u'(\hat{x}^v_*) = 0$. Here A_- and A_+ are independent of $\hat{x}$ but are undetermined, and may depend on both y and t. Define

$$\Phi_\pm = \frac{1}{A_\pm}\int_{\Omega_\pm} P_\pm\,\mathrm{d}\hat{x} \equiv \int_{\Omega_\pm}\phi\,\mathrm{d}\hat{x}. \tag{A15}$$

We note that $A_+\Phi_+ = 1-\theta$ introduced in section III F. These integrals can be evaluated using Laplace's method,

giving

$$\Phi_\pm \sim \sqrt{\frac{2\pi\delta}{\gamma_\pm}}\phi(\hat{x}^v_\pm), \tag{A16}$$

where $\hat{x}^v_\pm$ are the maxima of u (or the zeros of $\hat{v}$) in $\Omega_\pm$ and

$$\gamma_\pm = -\frac{\hat{v}'(\hat{x}^v_\pm)}{\hat{d}(\hat{x}^v_\pm)}.$$

In the one-dimensional case we could now use the normalisation condition on P to give a relationship between A_+ and A_-, namely $1 = \Phi_- A_- + \Phi_+ A_+$. But in the two-dimensional case we only have that

$$p(y,t) = \int_0^\infty P(\hat{x},y,t)\,\mathrm{d}\hat{x} = A_-(y,t)\Phi_- + A_+(y,t)\Phi_+. \tag{A17}$$

Define

$$p_\pm(y,t) = \Phi_\pm A_\pm(y,t), \tag{A18}$$

which we can interpret as the marginal density for Y given X is in the left/right well, multiplied by the probability that X is in that well. To calculate the exponentially slow transition rates we need to calculate the first eigenvalue/eigenfunction of $\mathcal{L}_\delta$, given by

$$\mathcal{L}_\delta \rho_1 = \lambda_1 \rho_1, \tag{A19}$$

say. Following Ward,[19] a good approximation of the first eigenfunction ρ_1 is the derivative of (A13) with respect to the unknown constant. Using $\rho_1(\hat{x}) = \partial P/\partial A_-$, gives

$$\rho_1(\hat{x}) \sim \begin{cases} \phi(\hat{x}) & \hat{x} < \hat{x}^v_*, \\ \frac{1}{2}\phi(\hat{x}^v_*)e^{-\gamma_*(\hat{x}-\hat{x}^v_*)^2/2\delta}\left[1 - \frac{\Phi_-}{\Phi_+} - \left(1+\frac{\Phi_-}{\Phi_+}\right)\mathrm{erf}\left(\sqrt{\frac{|\gamma_*|}{2\delta}}(\hat{x}-\hat{x}^v_*)\right)\right] & \hat{x} \approx \hat{x}^v_*, \\ -\frac{\Phi_-}{\Phi_+}\phi(\hat{x}) & \hat{x} > \hat{x}^v_*. \end{cases} \tag{A20}$$

To obtain λ_1 we use a spectral projection method that makes use of the adjoint operator $\mathcal{L}^*_\delta$, given by

$$\mathcal{L}^*_\delta \varphi \equiv \left(\delta\hat{d}(\hat{x})\frac{\partial^2\varphi}{\partial\hat{x}^2} + \hat{v}(\hat{x})\frac{\partial\varphi}{\partial\hat{x}}\right), \tag{A21}$$

together with boundary conditions $\varphi'(\hat{x}) = 0$ on $\hat{x} = 0, \infty$. The eigenfunctions of the adjoint operator satisfy

$$\mathcal{L}^*_\delta \xi_j = \lambda_j \xi_j, \tag{A22}$$

with the orthogonality relationship $\langle \rho_i, \xi_j\rangle = \delta_{ij}$ where δ_{ij} is the Kronecker delta. The adjoint eigenfunction corresponding to $\lambda_0 = 0$ is simply $\xi_0 = 1$. The first adjoint eigenfunction $\xi_1(\hat{x})$ is approximately (see e.g. Ref. 20)

$$\xi_1(\hat{x}) \sim \begin{cases} \frac{\Phi_+}{\Phi_-(\Phi_+ + \Phi_-)} & \hat{x} < \hat{x}^v_*, \\ \frac{\Phi_+}{\Phi_-(\Phi_+ + \Phi_-)} - \frac{1}{2\Phi_-}\mathrm{erf}\left(\sqrt{\frac{|\gamma_*|}{2\delta}}(\hat{x}-\hat{x}^v_*)\right) & \hat{x} \approx \hat{x}^v_*, \\ \frac{-1}{(\Phi_+ + \Phi_-)} & \hat{x} > \hat{x}^v_*, \end{cases} \tag{A23}$$

The first eigenvalue can now be computed by taking the inner product of (A19) with a suitable test function φ:

$$\langle \varphi, \mathcal{L}_\delta \rho_1\rangle = \lambda_1 \langle \varphi, \rho_1\rangle. \tag{A24}$$

Here we use $\varphi = \mathbb{1}_{\Omega_-}$. The left hand side gives

$$\begin{aligned} \langle \varphi, \mathcal{L}_\delta\rho_1\rangle &= \int_0^{\hat{x}^v_*} \mathcal{L}_\delta\rho_1\,\mathrm{d}\hat{x} = \left(\frac{\partial}{\partial\hat{x}}[\delta\hat{d}(\hat{x})\rho_1] - \hat{v}(\hat{x})\rho_1\right)\Big|_{\hat{x}=\hat{x}^v_*} \\ &= \sqrt{\delta}\hat{d}(\hat{x}^v_*)\rho_1'(z)|_{z=0} \\ &= -\sqrt{\delta}e^{u(\hat{x}^v_*)/\delta}\left(1+\frac{\Phi_-}{\Phi_+}\right)\sqrt{\frac{|\gamma_*|}{2\pi}}. \end{aligned} \tag{A25}$$

The right-hand side is

$$\langle \varphi, \rho_1\rangle = \int_0^{\hat{x}^v_*} \rho_1\,\mathrm{d}\hat{x} = \Phi_- \sim \sqrt{\frac{2\pi\delta}{\gamma_-}}\phi(\hat{x}^v_-), \tag{A26}$$

using (A16). Combining (A25) and (A26) gives the exponentially small eigenvalue

$$\lambda_1 \sim -\frac{\hat{d}(\hat{x}^v_-)}{2\pi}\left(1+\frac{\Phi_-}{\Phi_+}\right)\sqrt{\frac{|\gamma_*|}{\gamma_-}}\exp\left(\frac{u(\hat{x}^v_*) - u(\hat{x}^v_-)}{\delta}\right). \tag{A27}$$

Note that the argument of the exponential is negative since $u(\hat{x}^v_*) < u(\hat{x}^v_-)$.

Finally, we seek the differential equations describing the evolution of A_+ and A_-. First we integrate the equation (A10) with respect to $\hat{x}$ (equivalent to taking

the inner product with the adjoint eigenfunction $\xi_0 = 1$):

$$\hat{\epsilon}\left\langle 1, \frac{\partial P}{\partial t} + \frac{\partial J_y}{\partial y}\right\rangle = \langle \xi_0, \mathcal{L}_\delta P\rangle = \langle \mathcal{L}_\delta^* \xi_0, P\rangle = 0. \quad \text{(A28)}$$

From the left-hand side we have

$$\left\langle 1, \frac{\partial P}{\partial t}\right\rangle = \Phi_- \frac{\partial A_-}{\partial t} + \Phi_+ \frac{\partial A_+}{\partial t}, \quad \text{(A29)}$$

and

$$\begin{aligned}\left\langle 1, \frac{\partial J_y}{\partial y}\right\rangle &= \int_{\Omega_-} \frac{\partial}{\partial y}\left(V(\hat{x},y)A_- - \frac{\partial}{\partial y}[D(\hat{x},y)A_-]\right)\phi(\hat{x})\,\mathrm{d}\hat{x} \\ &\quad + \int_{\Omega_+} \frac{\partial}{\partial y}\left(V(\hat{x},y)A_+ - \frac{\partial}{\partial y}[D(\hat{x},y)A_+]\right)\phi(\hat{x})\,\mathrm{d}\hat{x}.\end{aligned} \quad \text{(A30)}$$

Taking the integrals over $\Omega_\pm$ inside the y derivatives and using that $\rho_\pm(\hat{x}) = \phi(\hat{x})/\Phi_\pm$, we find

$$\begin{aligned}\left\langle 1, \frac{\partial J_y}{\partial y}\right\rangle &= \Phi_- \frac{\partial}{\partial y}\left(\overline{V}_-(y)A_- - \frac{\partial[\overline{D}_-(y)A_-]}{\partial y}\right) \\ &\quad + \Phi_+ \frac{\partial}{\partial y}\left(\overline{V}_+(y)A_+ - \frac{\partial[\overline{D}_+(y)A_+]}{\partial y}\right),\end{aligned} \quad \text{(A31)}$$

where

$$\overline{F}_\pm(y) := \int_{\Omega_\pm} F(\hat{x},y)\rho_\pm(\hat{x})\,\mathrm{d}\hat{x}, \quad \text{for } F = D,\ V. \quad \text{(A32)}$$

Thus

$$\begin{aligned}0 &= \Phi_- \frac{\partial A_-}{\partial t} + \Phi_+ \frac{\partial A_+}{\partial t} \\ &\quad + \Phi_- \frac{\partial}{\partial y}\left(\overline{V}_-(y)A_- - \frac{\partial[\overline{D}_-(y)A_-]}{\partial y}\right) \\ &\quad + \Phi_+ \frac{\partial}{\partial y}\left(\overline{V}_+(y)A_+ - \frac{\partial[\overline{D}_+(y)A_+]}{\partial y}\right).\end{aligned} \quad \text{(A33)}$$

Next we take the inner product of (A10) with the first adjoint eigenfunction:

$$\hat{\epsilon}\left\langle \xi_1, \frac{\partial P}{\partial t} + \frac{\partial J_y}{\partial y}\right\rangle = \langle \xi_1, \mathcal{L}_\delta P\rangle = \langle \mathcal{L}_\delta^* \xi_1, P\rangle = \lambda_1 \langle \xi_1, P\rangle. \quad \text{(A34)}$$

Using (A13) and (A23) gives another PDE for A_- and A_+:

$$\begin{aligned}&\frac{\partial A_-}{\partial t} - \frac{\partial A_+}{\partial t} - \frac{\partial}{\partial y}\left(\frac{\partial[\overline{D}_-(y)A_-]}{\partial y} - \overline{V}_-(y)A_-\right) \\ &\quad + \frac{\partial}{\partial y}\left(\frac{\partial[\overline{D}_+(y)A_+]}{\partial y} - \overline{V}_+(y)A_+\right) \\ &\quad = \frac{\lambda_1}{\hat{\epsilon}}(A_- - A_+).\end{aligned} \quad \text{(A35)}$$

Rearranging (A33) and (A35) we find the following system for $A_-(y,t)$ and $A_+(y,t)$:

$$\begin{aligned}&\frac{\partial A_-}{\partial t} - \frac{\partial}{\partial y}\left(\frac{\partial[\overline{D}_-(y)A_-]}{\partial y} - \overline{V}_-(y)A_-\right) \\ &\quad = \frac{\lambda_1}{\hat{\epsilon}}\frac{\Phi_+}{\Phi_- + \Phi_+}(A_- - A_+),\end{aligned} \quad \text{(A36a)}$$

$$\begin{aligned}&\frac{\partial A_+}{\partial t} - \frac{\partial}{\partial y}\left(\frac{\partial[\overline{D}_+(y)A_+]}{\partial y} - \overline{V}_+(y)A_+\right) \\ &\quad = \frac{\lambda_1}{\hat{\epsilon}}\frac{\Phi_-}{\Phi_- + \Phi_+}(A_+ - A_-).\end{aligned} \quad \text{(A36b)}$$

Using (A18), we can write (A36) in terms of probabilities, $p_\pm(y,t)$:

$$\begin{aligned}&\frac{\partial p_-}{\partial t} - \frac{\partial}{\partial y}\left(\frac{\partial[\overline{D}_-(y)p_-]}{\partial y} - \overline{V}_-(y)p_-\right) \\ &\quad = \frac{\lambda_1}{\hat{\epsilon}}\frac{1}{\Phi_- + \Phi_+}(\Phi_+ p_- - \Phi_- p_+),\end{aligned} \quad \text{(A37a)}$$

$$\begin{aligned}&\frac{\partial p_+}{\partial t} - \frac{\partial}{\partial y}\left(\frac{\partial[\overline{D}_+(y)p_+]}{\partial y} - \overline{V}_+(y)p_+\right) \\ &\quad = \frac{\lambda_1}{\hat{\epsilon}}\frac{1}{\Phi_- + \Phi_+}(\Phi_- p_+ - \Phi_+ p_-).\end{aligned} \quad \text{(A37b)}$$

Finally, rescaling time in (A37) back to the original time variable [recall that in (A1) we had scaled time with τ_y], we find that

$$\frac{\partial p_-}{\partial t} - \frac{1}{\tau_y}\frac{\partial}{\partial y}\left\{\frac{\partial[\overline{D}_-(y)p_-]}{\partial y} - \overline{V}_-(y)p_-\right\} = k_+ p_+ - k_- p_-, \quad \text{(A38a)}$$

$$\frac{\partial p_+}{\partial t} - \frac{1}{\tau_y}\frac{\partial}{\partial y}\left\{\frac{\partial[\overline{D}_+(y)p_+]}{\partial y} - \overline{V}_+(y)p_+\right\} = k_- p_- - k_+ p_+, \quad \text{(A38b)}$$

where

$$k_- = \frac{\Phi_+}{\Phi_- + \Phi_+}\frac{|\lambda_1|}{\epsilon}, \qquad k_+ = \frac{\Phi_-}{\Phi_- + \Phi_+}\frac{|\lambda_1|}{\epsilon}. \quad \text{(A39)}$$

These rates can be identified as the transition rates from the left to right well and vice versa (introduced in section III F).

Equations (A38) are the Fokker–Planck equations associated with the reduced stochastic model (17a) appropriate for Regime 2 when Y is approximated by a continuous random variable.

3. Regime 3: $\epsilon \sim \tau_y \ll \tau_s$

In Regime 3 both X and Y switch between localised metastable states. A similar analysis to that in §A 2 can be used. In our simplified example in which the bistable variable X is independent of Y the switching rate is exactly given by §A 2; all that remains is to calculate

the quasi-stationary density for each metastable state. To approach this analytically requires a two-dimensional WKB (ray theory) approach, which is considerable more complicated than the one-dimensional version in §A 2.

If the system was fully coupled, a two-dimensional version of the eigenvalue calculation of §A 2 would be required to analytically determine the transition rates.[47] Alternatively we could numerically obtain the quasi-stationary densities of each well and the mean switching times between attractors using short bursts of stochastic simulation, as described in Sec. III F. Such a numerical approach could in principle be extended to an arbitrary number of attractors and/or higher dimensions as an automated process. However, as we have seen, care needs to be taken to define the boundaries between attractors and in determining the switching times.[22]

[1] P. Thomas, A. V. Straube, and R. Grima, BMC Syst. Biol. **6**, 39 (2012).
[2] H. H. McAdams and A. Arkin, Trends Genet. **15**, 65 (1999).
[3] E. M. Ozbudak, M. Thattai, H. N. Lim, B. I. Shraiman, and A. van Oudenaarden, Nature **427**, 737 (2004).
[4] J.-W. Veening, W. K. Smits, and O. P. Kuipers, Annu. Rev. Microbiol. **62**, 193 (2008).
[5] W. Horsthemke and R. Lefever, *Noise-induced transitions: Theory and applications in physics, chemistry, and biology*, 1st ed., Vol. 15 (Springer-Verlag, Berlin and New York, 1984).
[6] R. Wang, J. A. Dearing, P. G. Langdon, E. Zhang, X. Yang, V. Dakos, and M. Scheffer, Nature **492**, 419 (2012).
[7] W. E, D. Liu, and E. Vanden-Eijnden, Comm. Pure Appl. Math. **58**, 1544 (2005).
[8] A. C. Staver, S. Archibald, and S. A. Levin, Science **334**, 230 (2011).
[9] M. Scheffer, S. R. Carpenter, T. M. Lenton, J. Bascompte, W. Brock, V. Dakos, J. van de Koppel, I. A. van de Leemput, S. A. Levin, E. H. van Nes, M. Pascual, and J. Vandermeer, Science **338**, 344 (2012).
[10] N. Fenichel, J. Differ. Equations **31**, 53 (1979).
[11] M. Desroches, J. Guckenheimer, B. Krauskopf, C. Kuehn, H. M. Osinga, and M. Wechselberger, SIAM Rev. **54**, 211 (2012).
[12] P. Boxler, Probab. Theory Rel. **83**, 509 (1989).
[13] G. W. A. Constable, A. J. McKane, and T. Rogers, J. Phys. A: Math. Theor. **46**, 295002 (2013).
[14] G. Wainrib, Electronic Communications in Probability **18**, 1 (2013).
[15] C. W. Gardiner, *Handbook of Stochastic Methods for Physics, Chemistry and the Natural Sciences*, 3rd ed. (Springer-Verlag, New York, 2004).
[16] Y. Cao, D. T. Gillespie, and L. R. Petzold, J. Chem. Phys. **122**, 014116 (2005).
[17] C. V. Rao and A. P. Arkin, J. Chem. Phys. **118**, 4999 (2003).
[18] J. M. Newby and S. J. Chapman, J. Math. Biol. (2013).
[19] M. J. Ward, in *Analyzing Multiscale Phenomena Using Singular Perturbation Methods*, edited by J. Cronin and R. O'Malley (AMS publications, Providence, RI, 1998) pp. 151–184.
[20] R. Hinch and S. J. Chapman, Eur. J. Appl. Math **16**, 427 (2005).
[21] J. M. Newby, Phys. Biol. **9**, 026002 (2012).
[22] C. Hartmann, R. Banisch, M. Sarich, T. Badowski, and C. Schütte, Entropy **16**, 350 (2013).
[23] F. Schlögl, Z. Phys. **253**, 147 (1972).
[24] R. Erban, S. J. Chapman, I. G. Kevrekidis, and T. Vejchodský, SIAM J. Appl. Math. **70**, 984 (2009).
[25] Some authors would use the convention that the propensity function for the third reaction, $2X \xrightarrow{k_3/\epsilon} 3X$, is $k_3 x(x-1)/2\epsilon$ instead of $k_3 x(x-1)/\epsilon$. See Ref. 26 for a discussion on conventions regarding reaction rates.
[26] D. T. Gillespie, J. Phys. Chem. **81**, 2340 (1977).
[27] M. A. Gibson and J. Bruck, J. Phys. Chem. A **104**, 1876 (2000).
[28] Y. Cao, D. T. Gillespie, and L. R. Petzold, J. Chem. Phys. **124**, 044109 (2006).
[29] D. T. Gillespie, J. Chem. Phys. **113**, 297 (2000).
[30] This is in contrast to the Parallel Replica Algorithm,[32] which defines quasi-stationary distributions using absorbing boundary conditions.
[31] H. Risken, *The Fokker-Planck Equation*, Methods of Solution and Applications (Springer, 1996).
[32] C. Le Bris, T. Lelièvre, M. Luskin, and D. Perez, Monte Carlo Methods Appl. **18**, 119 (2012).
[33] O. Ovaskainen and B. Meerson, Trends Ecol. Evol. **25**, 643 (2010).
[34] C. E. Brassil, Ecol. model. **143**, 9 (2001).
[35] B. Dennis, Oikos **96**, 389 (2002).
[36] S. J. Schreiber, Theor. Popul. Biol. **64**, 201 (2003).
[37] G. M. Palamara, G. W. Delius, M. J. Smith, and O. L. Petchey, J. Theor. Biol. **334**, 61 (2013).
[38] I. Nåsell, J. Theor. Biol. **211**, 11 (2001).
[39] P. Turchin, *Complex Population Dynamics: A Theoretical/Empirical Synthesis* (Princeton University Press, 2003).
[40] J. Grasman and R. HilleRisLambers, Ecol. model. **103**, 71 (1997).
[41] This is because if we waited long enough in the stochastic model, extinction would eventually occur, and therefore the stationary density is not defined. In other words, there is a "leak" at $x = 0$.
[42] D. A. Kessler and N. M. Shnerb, J. Stat. Phys. **127**, 861 (2007).
[43] M. Assaf and B. Meerson, Phys. Rev. E **81**, 021116 (2010).
[44] J.-H. Prinz, H. Wu, M. Sarich, B. Keller, M. Senne, M. Held, J. D. Chodera, C. Schütte, and F. Noé, J. Chem. Phys. **134**, 174105 (2011).
[45] E. Abad, J. Reingruber, and M. S. P. Sansom, J. Chem. Phys. **130**, 085101 (2009).
[46] W. Chen, R. Erban, and S. J. Chapman, SIAM J. Appl. Math. **74**, 208 (2014).
[47] P. C. Bressloff and J. M. Newby, SIAM J. Appl. Dyn. Syst. **12**, 1394 (2013).
[48] V. Dakos, S. R. Carpenter, W. A. Brock, A. M. Ellison, V. Guttal, A. R. Ives, S. Kéfi, V. Livina, D. A. Seekell, E. H. van Nes, and M. Scheffer, PLoS ONE **7**, e41010 (2012).
[49] M. Geissbuehler and T. Lasser, Opt. Express **21**, 9862 (2013).
[50] P. C. Bressloff and J. M. Newby, in *First-Passage Phenomena and Their Applications*, edited by R. Metzler, G. Oshanin, and S. Redner (World Scientific, 2013) pp. 1–29.